

# NODES

## VOLUME II

.....

Stories of EEMCS









# EEMCS NODES



VOLUME II





# A treasure trove of stories

## Preface

Dear Lucas,

Over the past years, I've often been asked the question: *John, this faculty of yours... what actually goes on there?* I might be asked it during a New Year's reception or at a major international conference, or even at a family birthday party. So I'm really thrilled to be able to produce a book from my briefcase that describes all the fine research we do – beautifully written in simple language.

Just like the first edition, this book is a treasure trove of stories – stories that are set within our faculty walls but that reverberate far beyond. These are stories that generate a sense of optimism. For example, you can read how you can use mathematics to determine the germination capacity of potatoes and so contribute to addressing the ever growing global demand for food (*Health*), how computer graphics are set to broaden our experience of art (*Digital society*) and we will soon be able to tackle worldwide energy challenges in the new ESP Lab (*Energy transition*). What one word links all these EEMCS stories? Impact! Each and every one of these research projects puts something in motion. I am very proud of this. And I am very proud to present you with this great book.

All I can say is: read and be amazed. And at the same time, be aware that these are just a few of our stories. You will undoubtedly read more in a next edition in this series. Because packing all the EEMCS treasures into one volume would result in a book no letterbox in the Netherlands could accommodate. Lucas, I wish you a lot of success in the coming years and I'm sure you're going to enjoy all the wonderful research that is taking place within walking distance of your new office.

Have a good read!

**John Schmitz**

Dean of the Faculty of Electrical Engineering,  
Mathematics and Computer Science (2016-2020)

Hello John,

What a warm bath, and what a great start! In a world in which Covid-19 is striking mercilessly and indiscriminately, and many of us are forced to work from home, this book is a very welcome gift. What I would most like to do is to personally shake everyone's hand and listen to all the research project stories while enjoying a cup of coffee. But – sadly – that is not possible. That is why this book makes me even happier. It is the best and the quickest way to get to know the faculty. John, thank you so much for everything. In the coming years I will do my very, very best to continue all your good work.

Kindest regards,


**Lucas van Vliet**

Dean of the Faculty of Electrical Engineering,  
Mathematics and Computer Science (from 2021)

.....





# Table of contents



Nº	TITLE	PAGE
01	A fresh perspective on potato growth	6
02	Smart data collected from grandmother's living room	12
03	A bit of math to constrain epidemics	20
04	The impact of algorithms	26
05	Algorithms unravel cellular dynamics, and what it means for targeted therapy	34
06	Increased road safety with highresolution automotive radar	42
07	You only know the best method of defence when you understand more about your opponent	50
08	Protecting the electricity grid of the future	56
09	Computer as a sheepdog	60
10	Intermittent computing to replace trillions of batteries	66
11	Developing a digital twin for the electricity grid	72
12	How smart sensors can prevent epilepsy	76
13	Modelling a flexible electricity system	84
14	New Electrical Sustainable Powerlab smooths the way for energy transition	90
15	Understanding noise – from quantum fluctuations to climate models	96

Nº	TITLE	PAGE
16	Using mathematics to explain why the River Ems is silting up	102
17	From micro to macro with duality	108
18	Affordable MRI	114
19	Organizing high-dimensional data with triangles and fields	120
20	So much more than 3D visualisation	128
21	How not to waste energy on 5G	134
22	Is sharing personal data a matter of all or nothing?	142
23	Solar cells available at bargain prices, but still the revolution stagnates	146
24	A post-coronavirus society	152
25	Improving the effectiveness of speech recognition	156
26	Diagnosing tuberculosis with algorithms	164
27	Tackling tumours using mathematics	170
28	Ultrasound from a plaster	178
29	A glimpse into the universe with superconducting electronics	184
30	WeVisit: combatting loneliness with video-calling	192





## A fresh perspective on potato growth

..... Text: Merel Engelsman | Photography: Mark Prins

Requiring less water and less time to grow than rice, corn or wheat, potatoes may well be the answer to the growing worldwide demand for food. But what are the key factors that determine a successful potato harvest? Asked to provide a fresh perspective, Neil Budko develops numerical methods for the analysis of carefully controlled potato growth experiments.

01



*'It's a typical inverse problem - the observed vitality is the effect and now we want to determine its causes.'*

### A bit of potato fun

'About ten years ago, I switched from electrical engineering to my current position,' says Neil Budko, associate professor in the Numerical Analysis Group at TU Delft. 'It was a shift towards research with a more direct practical application. I like that.' His first exposure to agriculture was a quirky little assignment from a Dutch company specialized in seed potatoes. 'We met at a yearly industrial-academic workshop where companies pose problems to a group of mathematicians, to be solved within a week. They wanted to know how to best cut potatoes into fries,' he says. 'It was a fun project, although scientifically not very challenging. After a few more of such projects, however, I realized there is something in seed potatoes that fits well with my background in inverse problems. That is a class of problems in which the outcome is known and where you want to determine the causal factors that produced this outcome.'

### Ideal potatoes don't grow

Potatoes are so-called *stem tubers*. A tuber is a storage organ for nutrients, grown by certain plant species to survive the winter or dry spells. Potato plants grow up to twenty of such underground tubers, to be used for consumption or for growing new plants. In the latter case, such a tuber is called a seed potato and it is planted at a depth

of about twelve centimetres. It develops sprouts, which grow into stems that will eventually emerge from the ground. These stems develop leaves. The more (large) leaves they have, the more excess nutrients to be stored and the more potatoes to be harvested. 'Many factors contribute to potato growth,' Budko says. 'The potato variety and its genetic constitution, for example, but also the weather, soil, growth and storage history of the seed tubers.' Agronomists have spent decades building a model for growing an ideal potato, resulting in a very complicated equation. 'This model provides a lot of understanding,' Budko says, 'but it doesn't really work as there are no ideal potatoes. We were contacted by two seed potato companies to provide a fresh perspective on potato growth.'

Neil Budko, associate professor in the Numerical Analysis Group at TU Delft.

### The Flight to Vitality project

The soil in the north of the Netherlands is especially well-suited for growing seed potatoes. Farmers buy these seeds from the producers in order to grow potatoes for actual consumption. Both producers and farmers would like to know the vitality of seed potatoes prior to planting. 'Vitality is defined as the proportion of seed potatoes that emerges from the soil and their uniformity of growth after emergence,' Budko explains. 'The current standard for determining their vitality is to pick a bunch of potatoes, put them in a bucket for a week, and see how many become rotten. You would think this methodology is easy to beat, but it is actually very hard to come up with a systematic procedure.' This is what the Flight to Vitality project is about, which also involves the two seed potato companies and Utrecht University. 'As the major contributor, we are responsible for all the math and data processing during the project.'

### Climate control and drones

'At the end of this project, we want to be able to just look at a batch of seed potatoes, measure the few aspects that we figured out to be important, and exactly tell you the batch's vitality,' Budko says. To determine these important aspects, the researchers designed a very broad search, focussing on the performance of six potato varieties in all kinds of conditions. Tubers were simultaneously planted in various countries and soils as well as in two custom-build climate-control rooms – a cold and a warm room, each having a dry and a wet section. For each batch of tubers, a lot of data are collected and stored: genetics, microbiomics, storage conditions, organic



*'Ideal models don't work. They came to us for a fresh perspective on potato vitality.'*





and inorganic constitution. Once planted detailed weather data are collected as well. Camera-equipped drones are used to monitor plant growth – when do they first appear above ground and how do they continue to develop. Acquisition of images at wavelengths outside the spectrum of visible light allows the researchers to also track changes in the biochemical composition of the plants. ‘It’s a typical inverse problem,’ Budko says. ‘We observe the measured vitality, and now we want to determine what caused these effects. We get thousands of signatures, each consisting of thousands of data points. Gigantic data matrices.’

### Back to natural intelligence

‘We can’t use an artificial neural network to analyse these data, as we do not have a classification problem but a so-called *regression problem*,’ Budko explains. ‘Artificial neural networks are extremely good at classifying things – pictures of cats and dogs, or CT-scans with or without cancer – and bad at telling you how they do that. We want to understand the relationships in the data, to determine the root causes of potato vitality. We will then perform additional experiments to validate the cause-effect relationships we have found. Further-

more, setting up and conducting experiments and collecting data is very expensive and we may never have a training set large enough for automated machine learning.’

### Leaf it to the computer

Perhaps the most important aspect of the ongoing research is to measure the vitality of each batch of seed potatoes, in which Budko is assisted by his PhD Student Elisa Atza and his master’s student Martijn Bos. Taking drone images is standard industry practice. But rather than a one-time snapshot, the researchers want detailed growth information over a period of four to eight weeks. As each photo only covers a small section of the potato fields, they first need to be stitched together. They then need to align these aggregate pictures in space and in time, and also extract the growth information. ‘We are developing an approach to automatically determine the change between pictures,’ Budko says. ‘Not just the difference in the pixels, but the actual change in the size of the plants, in each of the leaves. How much have they grown in length, how much in width? Existing algorithms for such deformable image registration are not subtle enough. We had to develop our own algorithms, so as not

to transform away any of the growth.’ The gigantic transformation process, involving hundreds of pictures each day and thousands of plants per picture, has to be automated. ‘Our computers run out of memory. We have some tricks up our sleeves, but it is not solved yet.’

### Yielding control, controlling yield

‘Obviously, it will be the seed potato’s history and its chemical constitution that influence vitality,’ Budko says. ‘What we don’t know yet, however, is the relative importance of each of these factors. The exact proportions of specific organic and inorganic substances, for example. But potato vitality is starting to expose some of its secrets to us.’ Vitality and yield are related. If a plant has emerged and temperature and humidity are okay, it will inevitably produce potatoes. ‘That is why we focus on the first few weeks of plant growth,’ says Budko. ‘Maximizing production is not the goal of our project. But if we know what determines yield, and it is easy to control, why not?’

### Watching potatoes grow in real time

The Flight to Vitality project was already under way when Budko became aware of another longstanding problem in the root vegetables business. As all root vegetables grow underground, they need to be dug out for inspection – a wasteful and time-consuming procedure. ‘It reminded me of a project I worked on as an electrical engineer – detecting landmines with ground radar. I also remembered the limitations of that technique, so we had to come up with some other solution.’ Air temperature varies, with warmer temperatures during the day and colder temperatures at night. Budko and his bachelor student Marissa Bezemer modelled how this affects the temperature in the soil and saw a temperature wave moving down. They then added the potatoes to this model. ‘In soup, potatoes are always the hot bastards,’ Budko says. ‘They have a very high heat capacity, which leads to local distortions in the temperature wave. We ran the simulation for potatoes of various sizes and got a very cool movie.’ Budko now wants to put an array of temperature sensors into the ground, leaving it there for the whole growth season. ‘The company involved is very excited, as nobody can yet measure potato growth non-destructively.’

*‘The current standard for determining potato vitality is to put them in a bucket for a week and see how many become rotten.’*






## 02

# Smart data collected from grandmother's living room

..... Text: Robert Visscher | Photography: Mark Prins

Chatting with family on the other side of the world via a hologram or selecting your own camera angles for your favourite club's football match. TU Delft computer scientist Pablo Cesar is researching the best ways of transmitting and receiving data to make these kinds of applications possible. He was recently awarded the Netherlands Prize for ICT Research 2020.



At first glance, the room where Pablo Cesar works looks like a totally normal office complex, with carpeting, system ceilings and a pinboard with posters about academic events. Behind the doors, you expect to see desks with computers on which mathematicians are doing calculations. But when Cesar opens the door, you walk into a cosy traditional Dutch living room, one your grandmother would feel at home in. There is a photo of the king next to a cuckoo clock, and a Persian rug on the floor.

But this is not where the scientists drink coffee or relax watching television. 'This is one of our research labs,' says Cesar. He points towards an antique lamp. 'It is fitted with a microphone and there are cameras on the walls, so we can analyse user behaviour. Next door is a computer room where we test and compare different TV and data configurations.'


*'In order to understand human behaviour in the wild, such as dancing in a club, we used machines with artificial intelligence quickly processing data from sensors from around 450 people simultaneously.'*





About Pablo

**Pablo Cesar** (1976) has been appointed professor of Human-Centered Multimedia Systems in the Department of Intelligent Systems (INSY) at TU Delft (Faculty of Electrical Engineering, Mathematics and Computer Science). Cesar also leads the Distributed and Interactive Systems (DIS) group at CWI and is the recipient of the Netherlands Prize for ICT Research 2020. His research focuses on modelling and controlling complex collections of media objects (including real-time media and sensor data) that are distributed in time and space. In his research, Cesar in particular focuses on fundamental problems related to the interaction between people and technology, and how this interaction can be modelled. The research of Cesar is not restricted to the lab, but his experiments as well performed in a realistic context.



.....

Our research group focuses on facilitating and improving the way people access media and communicate with others and with the environment. We address key problems for society and science, resulting from the dense connectivity of content, people, and devices. We use recognized scientific methods, following a full-stack, experimental, and human-centered approach.

*‘We prefer to conduct our research in the wild.’*

Prestigious prize

The lab disguised as grandmother’s living room provides an indication of how Cesar and his colleagues approach their work. They like to conduct research in as realistic a setting as possible. They use this method to look at a wide range of subjects. For example, Cesar is investigating how we watch TV, he aims to improve teleconferencing and wants to enable hospitals and ill people to communicate with each other.

The computer scientist is an associate professor in TU Delft’s Multimedia Computing group and the national research institute for mathematics and computer science Centrum Wiskunde & Informatica (CWI) in Amsterdam. For his contribution to science, he was recently awarded the Netherlands Prize for ICT Research 2020. In its report, the judging panel highlighted the fact that Cesar’s work focuses on fundamental problems in human/technology interaction and how these can be modelled. They also praised Cesar for conducting his experiments in a realistic context, rather than a lab setting. ‘We prefer to conduct our research in the wild,’ says Cesar. ‘We’re developing technology for people. That means that we have to get them to test what we come up with in order to improve our ideas.’



*'It's normally the director who determines what you see at home. I am interested in exploring mechanisms for content customisation and viewer personalisation.'*

A good example of his approach is the British FA Cup final. Cesar was asked by the BBC and British Telecom to design a second screen for people watching the final. He came up with a menu that allows you to control all kinds of things, including the camera angle, extra action replays or the colours around the screen. 'It's normally the programme director that determines what you see at home. Our technology gives the viewer more freedom. The director is still in control, as is appropriate in view of his or her experience. But he or she gives the viewer multiple options to choose from.'

This made it possible to show different things on several screens in a pub full of Chelsea supporters. One TV showed images of the match whereas another just had a close-up of the coach. The sound came directly from the Chelsea supporters' stand, so their chanting could be heard in the pub. The result was much more tailored to the specific target group in the pub.

This approach may give viewers more options, but what difference does the science make here? 'Our primary interest was in the how to transmit all that information at a superfast speed and orchestrate the complex data on the screens. Pictures of the match already need to be transmitted to the TV in the pub or at home at amazing speed. In our case, we had to send and control a lot more data, such as several camera angles and different sound.'

That involves transmitting an awful lot of information simultaneously, because you also want to avoid any delay in the match you're watching. Cesar is developing ways of transmitting and controlling the data as smartly as possible. To achieve this, he uses algorithms that select the relevant information in a smart way. This data only is all put in small packets in a process known as compression – almost like creating a zip file on your computer. The files are pressed together, reducing the space they take up and enabling them to be transmitted faster and more easily. The files are then unpacked at the end user where they are converted to make images or sound. Algorithms also play an important role here. Cesar is investigating how all of this can be done as effectively and quickly as possible. In the process, he always has the viewer's quality of experience in mind. What compression results in the best image quality, how much image delay is permitted without undermining the connection and what bandwidth is required?

### Communicating with family

Cesar is not only working on watching TV and second screens. He's also researching communication between people. He walks into a room full of boxes. They have been painted pink on one side. There are cameras around them and a researcher is walking round with his arms spread out. Cesar points to a computer screen showing an image of the walking man.



Using the technology that Cesar is now working on, you can create in real-time a highly realistic 3D representation of someone that makes it look as if he or she is standing in front of you.

Imagine you want to discuss something important with a family member on the other side of the world. It's now already possible to telephone them or use Skype. But using the technology that Cesar is now working on, you can create a hologram of someone that makes it look as if he or she is standing in front of you. He is working on this in this room.

'We create a 3D image of someone. The result is a large file that we compress and decompress when it reaches the receiver. It's similar to what we did for the FA Cup. But what's the most practical approach? If I want to talk to a family member in Spain, where I come from, they don't need to see my back. That means that we don't have to send all of the images. These are very large files, so we're investigating how we can make them smaller by removing useless noise.'

The image quality required also plays a role. Displaying someone's face in high definition and the rest of the body less clearly can already make a big difference. Cesar and his colleagues are testing people's responses to this. 'We are trying to figure out which bits of data are important to send and receive and which are not.'

This up-and-coming technology is not only useful for phoning your family. It can also be used for teleconferences with businesses, or in hospitals. 'For example, if you need a knee operation, you can have the necessary appointments for tests in the hospital. Any other appointments can be done by a teleconnection, with the doctor receiving a very high definition image of the knee. This will also help the advance of healthcare in the future.'





## Disco of the future

But it's not only healthcare – even a disco can benefit from Cesar's research. Several years ago, Cesar was asked to work on a project at the Amsterdam Dance Event (ADE). One of the sponsors asked him to make the club of the future possible. Cesar therefore joined forces with clothing designer Borre Akkersdijk to come up with armbands packed with sensors that checked people's temperature, location and movements. This information was then transmitted wirelessly to minicomputers on the wall. Some 450 people attended. 'That's an amazing number for a scientific study,' says Cesar. 'We used machines with artificial intelligence that had to understand where people were and whether they were dancing. We then wanted to share that information with the people at the event. For that, we created a second room that was the opposite of the dance floor.'

When the feet really hit the floor, it was very relaxed in the second room. If everyone on the dance floor was acting like a wallflower, the second room was incredibly hectic. 'By doing this, we learnt how to process a lot of information rapidly and also how to display it immediately to visitors. The result was not only a great disco, but something that we, as scientists, could also learn from.' In scientific research into what is known as affective computing, there is a lot of focus on such issues as the recognition of emotions by software and investigating audience engagement.


In these kinds of projects, the researchers make a lot of things themselves. This not only includes the software, such as algorithms to do the calculations, but also the sensors built into the armband. Cesar reveals a room full of equipment where there is also a workbench on which all of the hardware is made or adapted.

This is not pure chance. Cesar and his colleagues are interested in the entire chain across which information is transmitted and received. It's an approach he believes in. 'We're living in an age when many scientists are focusing on a very small thing or subject. We look at things from a broader perspective and always focus specifically on how information reaches people. This has been relatively ignored in the scientific world, but our work at TU Delft and CWI is fortunately now changing that. Winning the ICT prize in 2020 will certainly help. I hope that the attention it attracts will focus especially on the development of computer science aimed at end users and that we will be able to continue to investigate that even more effectively.'



Cesar joined forces with clothing designer Borre Akkersdijk to come up with armbands packed with sensors that monitored people's temperature, location and movements, and helped to understand human behaviour.





## 03

# A bit of math to constrain epidemics

..... Text: Merel Engelsman | Photography: Mark Prins

Most people like to keep both mathematics and epidemics at a distance. But a little bit of the former can do wonders in limiting the latter. Supervised by Professor Kees Vuik, two high school students modelled a recent outbreak of cholera in Ghana. He also has a bachelor student investigating the new corona virus.

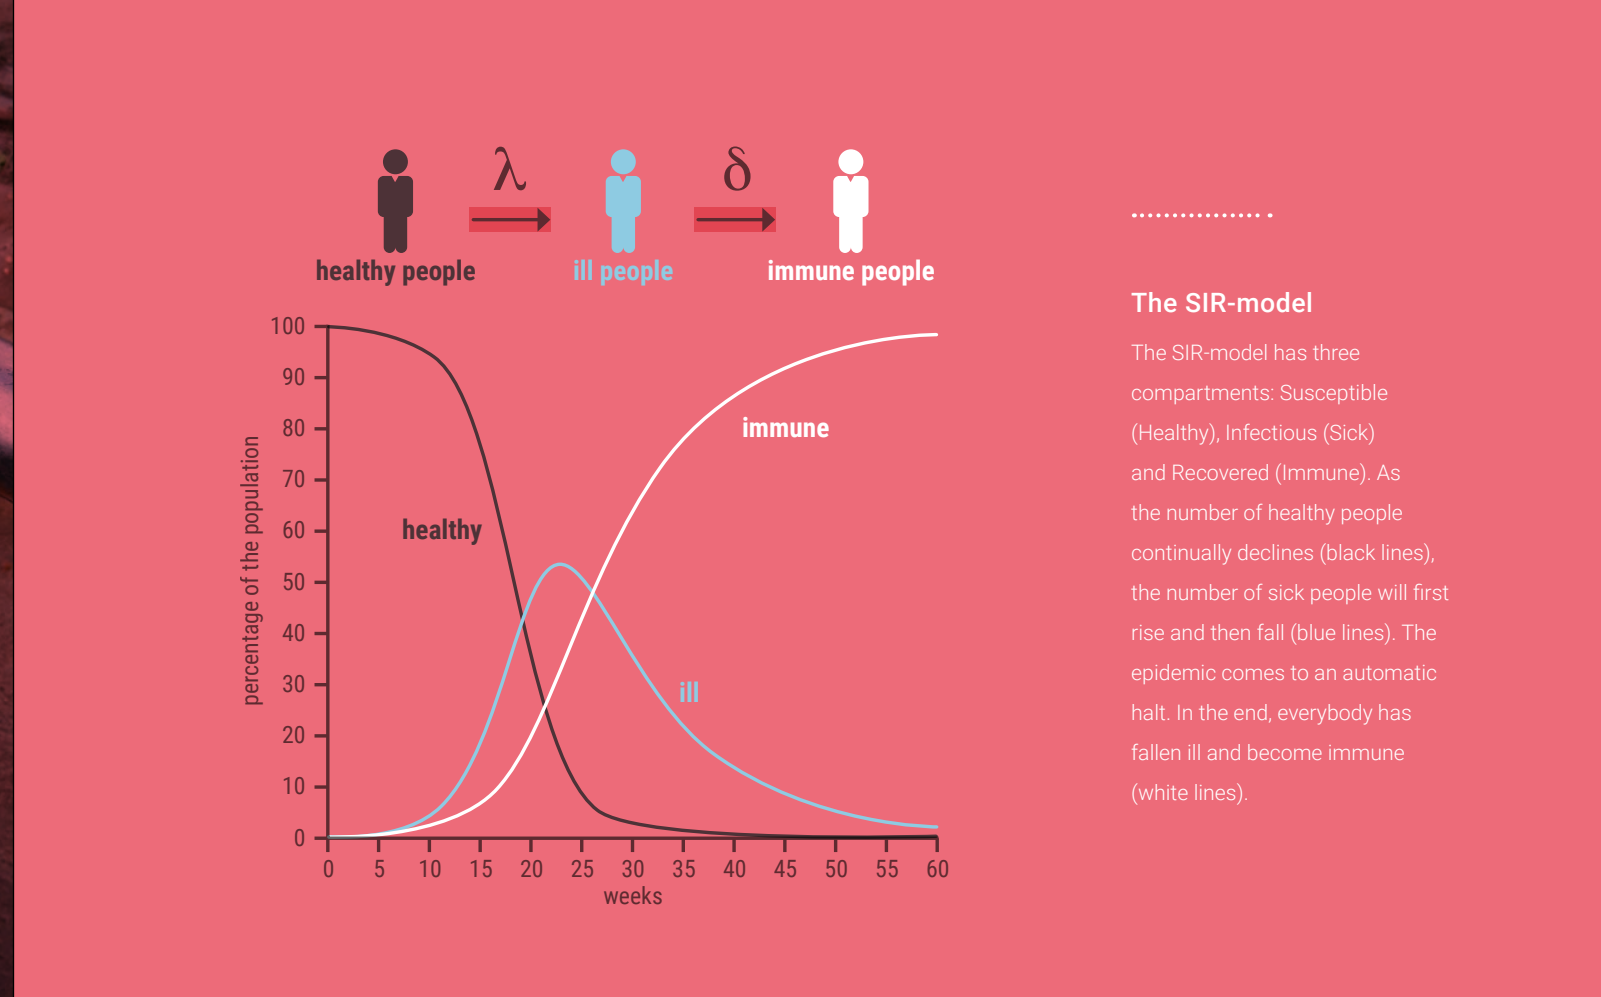
‘For ten years I have used the course of small outbreaks and large epidemics to introduce the concept of mathematical modelling to new university students,’ says Kees Vuik, professor in the Numerical Analysis group at TU Delft. ‘Although we do explore some mathematical depths, these models are essentially relatively straightforward. It would be great if everybody gained some insight into these models. We would then better understand why certain measures are, or are not, taken in an effort to control the course of an epidemic.’

### The dynamics of a simple epidemic

For well-known diseases such as the flu, measles, mumps and rubella, the course of an epidemic can be accurately predicted using the SIR-model (see figure). It requires the population to be divided into three groups, so-called *compartments*: Healthy people, Sick people and Immune people. Those in the group of Healthy people have a certain chance to fall ill, and thereby move to the group of Sick people. The chance of falling ill increases as the number of people that are already sick and therefore contagious increases. Those in the group of Sick people have a certain chance to recover and to thereby move to the group of Immune people. This very straightforward model correctly predicts there to be a peak in the number of sick individuals. The magnitude and timing

*‘With a bit more insight into disease modelling we would all better understand why certain measures are taken to constrain an epidemic.’*





*‘We wanted to do a project involving mathematics and we wanted it to have social relevance.’*

of this peak depends on two coefficients. The first coefficient is a number, the value of which signifies the speed of sick people infecting healthy people. The value of the second coefficient signifies the speed of recovery of sick people. The model can be extended with aspects such as vaccinations or a latent period during which a person is infected but not yet contagious. It is harder to mathematically solve such an extended model.

### Cholera in Ghana

‘Early last year, at a conference, I saw a poster describing a disease model for cholera,’ Vuik says. ‘The model showed two pathways for the cholera bacterium to be transmitted – from one human directly to another, and through

contaminated drinking water. That was new to me. I then asked two high school students to try and apply this model to a recent outbreak of cholera in Ghana.’ In 2014, more than seventeen thousand people were infected by cholera, with about three hundred of those not surviving. The two students, Annelotte Baarslag and Mark Tabaksblat of Gymnasium Novum in Voorburg, had just been accepted to the Pre-University College of Leiden University and were eager to get their teeth into it. ‘We had many assignments to choose from for our final research project. We wanted to do a project involving mathematics and we wanted it to have social relevance. Modelling a disease outbreak was very appealing to us.’

### No longer contained

‘We had to learn some additional math skills,’ Mark says, ‘but we could still explain our research and the model we used to our classmates and family. Not so much how to solve the model using numerical analysis, but certainly the various dependencies in the model and how they influence each other as well as the course of the outbreak.’ The students first gained experience by applying the simple SIR-model to an outbreak of influenza. They subsequently focussed their attention on the more difficult model for cholera, including the additional infection pathway through contaminated drinking water. The model was also extended by including both birth and natural death. Next to assuring that a model correctly represents all relevant dependencies, another big challenge is to determine the correct values of all its coefficients. ‘At first, Mark and Annelotte were a bit hesitant to as much as change the value

of a single coefficient,’ Vuik explains. ‘But as soon as they saw the predicted course of the epidemic to much better match the measured rate of infections and deaths, they could no longer be contained. They doubted the correctness of all coefficients and they even changed a dependency in the model.’

### A sewage system but no quarantine

‘With our model matching the outbreak in Ghana we could investigate the measures available to the authorities to constrain an outbreak of cholera,’ Annelotte says. ‘We for example lowered the value of the coefficient representing the speed with which the cholera bacteria transfers from humans to local waters. The authorities can achieve such a reduction by building a sewage system or by improving an existing one. Our model predicted this measure to reduce the total number of people being infected as well as the number of deaths caused by the outbreak.’





A futuristic scenario with extended insight  
'It is the year 2041. Following a first limited outbreak under high school students in Voorburg, mathematical understanding of disease models has spread all over the Netherlands. The entire city of Rotterdam is put under quarantine as, after nearly 200 years, it is unexpectedly hit by an outbreak of cholera in its port area. Nobody was surprised by this drastic measure and no parliamentary questions were tabled.'

This was an expected outcome as, mid-19th century, the Netherlands halted cholera outbreaks by adding a sewage system and a pipeline network for drinking water to many of its cities. 'Perhaps the most interesting outcome was that imposing quarantine measures has absolutely no effect in the case of cholera,' Mark says. 'Although it reduces the rate of infection between humans, everybody still ends up being infected through

contaminated drinking water. Quarantine or not, the same number of people will succumb to cholera.' Professor Vuik adds some nuance to this conclusion. 'A quarantine does work if you do not only quarantine people, but also the contaminated water sources. This can be achieved by restricting access to an entire village or region. More comprehensive disease models therefore subdivide an area or country into various regions, while also modelling travel between these regions.'

### It's in the details

'The model we used to describe the cholera outbreak is quite advanced and at a scientific level,' Vuik says, 'but the real know-how is concentrated at the National Institute for Public Health and the Environment (RIVM).' The institute houses twenty mathematicians dedicated to disease modelling. At this moment, with the corona outbreak, they work around the clock. 'It is a new disease for which the model dependencies are well-

*'And just to think that, at this moment, mathematicians are actually saving lives.'*

known, but the epidemic has only just started,' Vuik continues. 'Each day brings new data, leading to changed insights and updated values of coefficients: contagiousness, the duration of the latent period, and so on – per age category.' The models used by the RIVM are fundamentally similar to the students' cholera model, but they are extremely hard to solve mathematically as they need to take so many details into account. The simple models assume every person to have a similar likelihood to be infected, and to be equally contagious once infected. The model of the RIVM applies different coefficients per age group and gender. They also take into account highly detailed information about the interactions between people – who meets whom, how often and where? It allows the RIVM to determine the effectiveness of countless and quite detailed measures, such as the closing of schools. 'I currently also supervise a bachelor student. He is investigating whether or not some of these measures also follow from a simpler model (see frame).

### The power of mathematics

Both high school students were already familiar with the power of mathematics prior to their modelling assignment. It is something they wish to continue to apply. Annelotte has been admitted to the navy and will be trained to become a maritime officer – a bachelor involving a lot of mathematics and physics. Mark will most likely enrol at Leiden University to study mathematics and physics. 'The current corona epidemic is so overwhelming,' he says. 'And just to think that, at this moment, mathematicians are actually saving lives.'

SCAN!



Who is that chalk-spattered individual patiently explaining differential equations? Who is that man or woman struggling with the smartboard? And while we're at it... where did the unbridled passion for blockchain originate? These are just three of the questions you will get answers to from listening to De ProfCast, a EEMCS podcast created by PhD candidate Marieke Kootte and Master's student Tijs Ziere. The first episode will feature Kees Vuik (Numerical Analysis). Listen to the podcast (which is in Dutch by the way) here.

### A simple disease model for corona

A bachelor student supervised by Professor Vuik is currently investigating if a simpler disease model can (still) be adequate for understanding why certain measures are proposed to constrain the current corona epidemic, such as closing our airspace to flights from certain high-risk countries. The student will specifically try to explain how this virus was able to spread around the world so quickly. To answer this question, he will model the influence of long-distance travel in combination with the probability of being contagious while not yet experiencing any symptoms. 'I think it may be interesting to add a fourth category to the SIR-model, a contagious-but-not-yet-ill compartment,' says Professor Vuik. 'The most important aspect of this research is educating students. It is not our aim, or our expectation, to discover something that is of interest to the RIVM, but it is certainly a possibility.'



# The impact of algorithms

Text: Merel Engelsman | Photography: Mark Prins, Inge Hoogland

In our age of technology algorithms have become more complex and more influential than ever before. And yet, because of their omnipresence, we are much less aware of them. We only notice their existence when they stop doing what they were intended for. What exactly is an algorithm, and what is their impact?

## What is an algorithm?

Many people associate algorithms with a computer solving a complex mathematical problem. But in its essence, an algorithm is little more than a finite set of instructions for achieving a certain goal. A recipe for preparing a tasty meal also fits this description, but the term algorithm will indeed more often be used in combination with numbers and computers.

In general, these algorithms process information. Mathematics are important, for example for translating a real-world problem into a model, which then determines the so-called *solution space*. The algorithm is a smart way to search for the best possible solution within this space. A good algorithm is characterised by it being both correct and efficient. Efficiency means that the solution will be reached with as few steps – as little computing power – as possible. Thanks to an algorithm, the user does not need to understand the problem in order to solve it. The inventor of the algorithm, on the other hand, needs to have a thorough understanding on the subject.

*'Thanks to an algorithm, the user does not need to understand the problem in order to solve it.'*

## Euclid's algorithm

One of the oldest known (mathematical) algorithms, defined in 300 BC by Euclid, finds the greatest common divisor of two numbers. The algorithm shows great understanding by its inventor, is easy for the user to apply, and works as follows: Subtract the smallest number from the biggest number. Repeat for the new number and the smallest of the two previous numbers until you reach 0 (zero).

For the numbers 49 and 21, the steps are:

$$49 - 21 = 28, 28 - 21 = 7, 21 - 7 = 14,$$

$$14 - 7 = 7, 7 - 7 = 0.$$

The greatest common divisor of 49 and 21 is the number 7.








## The role of computers

A computer program is the implementation of an algorithm. Multiple computer programs can lead to the same solution, but they may differ in, for example, the number of calculations needed, or the amount of storage space. In the early days of the computer age, computational power was very much limited, and the science of algorithms focussed on the discovery and implementation of algorithms for problems that were relatively easy to solve, such as finding the shortest route. Because of the enormous increase in computing power and in the fundamental knowledge of algorithms, this focus has shifted towards problems that were previously unsolvable because of their sheer magnitude. Where the need for fast execution of algorithms used to be a driving force behind the development of computers, nowadays powerful computers and the availability of huge digital datasets have led to the development of new algorithms. Machine Learning is currently drawing most of the attention.

## Machine Learning

In Machine Learning, a kind of artificial intelligence, it is no longer human understanding that powers the algorithm. Rather, it is the algorithm itself that uses a large dataset to determine a relationship between input and output. In this process, the role of humans is limited to gathering a correct and sizable dataset for training the algorithm, designing the underlying artificial neural network (in case of the highly popular Deep Learning approach), and the decision of whether or not to release the trained network into real-world operation. Machine Learning has led to breakthrough successes in bioinformatics and in the fields of speech and image recognition. Machine Learning networks outperform humans in certain tasks, but they do not think like humans. In fact, their inner workings are so complex that it is often impossible for a human to determine how and why a Machine Learning network makes certain decisions. Their elusiveness may give them a somewhat spooky quality, as their influence on our daily life is continually increasing.

As a consumer you often don't realise when an algorithm impacts your life. When the train schedule is adapted because of a disruption, for example.



In most recent projects Peter Bosman focuses on topics in radiation oncology, including automated treatment planning, deformable image registration and 3D dose reconstruction, in close collaboration with researchers and medical experts from (academic) hospitals.

## Finite applicability

As a consumer you often don't realise when an algorithm impacts your life. When your phone automatically adjusts its brightness setting, for example, but also when the train schedule is adapted because of a disruption. In the latter case, an algorithm could perhaps instruct trains not to stop at your tiny train station for a few hours. It is therefore important to realise the circumstances under which an algorithm stops performing as intended. For certain algorithms, such as the one by Euclid (see page 27), it can be proven mathematically that they will always function correctly. For other algorithms, such as the one adjusting the train schedule, this may only be the case as long as the underlying model accurately represents reality. Removal of a train track may necessitate the use of a different algorithm, or at least a new optimisation using the existing algorithm. Deep Learning neural networks, and virtually all kinds of artificial intelligence, can find very good solutions, but it is not possible to mathematically prove these to be optimal. Their training datasets may furthermore be biased (all training pictures containing cats also contain furniture), resulting in unexpected network behaviour.

## New insights

The use of existing (combinations of) algorithms to solve real-world problems almost always results in suboptimal solutions that are either too slow or not accurate enough. To fix this, the underlying model may have to be over-simplified. A researcher will then implement smart improvements to build a suitable algorithmic solution. These algorithmic insights often have a significance beyond the scope of the problem just solved.

*'An optimal solution that isn't fair, is not optimal to me.'*

Mathijs de Weerd





Mathijs de Weerd

A fair energy transition

As associate professor in the Algorithmics group at TU Delft, Mathijs de Weerd employs algorithms for solving problems involving multiple parties. In his current research, he develops algorithms for ensuring an efficient and fair use of the electricity network. This is of special importance now that supply and demand are less and less in sync due to the transition towards renewable energy. 'You cannot generate solar power at night, or wind power on a calm day,' he says. 'It helps if we only charge our electric car and turn on our electric boiler at moments of sufficient supply.' Furthermore, there are capacity constraints within the electricity network. In theory, this can all be addressed by offering a financial incentive to parties that would like to supply or use electricity at peak moments. The magnitude of the incentive depends on how strong a need each party expresses. 'This only works when the optimal solution is found, each time,' explains de Weerd. 'Otherwise, users can abuse the system by expressing a fake need.' As an optimal solution is required, de Weerd does not use Machine Learning. He also insists for the solution to be a fair one. 'It is not fair to always cut off the same neighbourhood or windmill park in

case of capacity issues. This may be the best solution for all users as a group, but for me it is not an optimal solution.' In the long run, because of his algorithms, fewer gas and coal-fired power stations may be needed to handle network imbalances. His research is of both practical and scientific value. 'When modelling a real-world problem, we often encounter new challenges, such as a fair distribution. We then take the structure of these problems into account to build faster and better solutions,' he explains. 'This way, we build new algorithms containing components of existing algorithms. We then show that our new approach may be more widely applicable.'



*'Thanks to increased computing power, we can use much more detail when optimising the spread of emergency ambulances.'*

Theresia van Essen

Theresia van Essen

Optimal spread of ambulance services

With many medical professionals in her family, Theresia van Essen's childhood dream was to become a surgeon. But it is algorithmics with which she now helps to achieve optimal medical care. As assistant professor in the Optimization group at TU Delft she for example investigates the optimal spread of emergency care ambulance services. 'Healthcare is expensive, so any improvement in services, without increasing cost, is more than welcome,' Van Essen says. Because of increased computer power, this problem can now be addressed at a much more detailed level. 'In the Netherlands, our level of detail is down to the numeric part of the postal codes. We take into account the current spread of emergency ambulances and their historic response times, for example.' The starting point of her research is to find an optimal solution for the highest level of detail. 'If determining the optimal

solution takes too much time, I work on algorithms approaching that optimal solution in much less time. One could also simplify the model, of course, for example by assuming average response times that are constant throughout the day. But an optimal solution for a simplified model might not be optimal for the real problem.' Thanks to her algorithms, an additional dynamic ambulance post has been introduced in the Utrecht area, for a number of hours each day. As a result, the region now complies with the applicable regulations regarding the maximum response time. It is unlikely for the additional ambulance post to be cancelled anytime soon. 'The number of incidents in a region will only vary slowly over the course of years,' Van Essen explains. 'It is, however, still advisable to rerun the algorithm each year, to see if a substantial improvement in emergency ambulance care can be achieved after all.'



*'Our aim is to stimulate diverse behaviour in the user.'*

Emily Sullivan



Emily Sullivan

Diversity in recommendations

Emily Sullivan is a postdoc in the Web Information Systems group at the TU Delft, where she works on improving so-called *recommender systems*. These systems generate recommendations, based on the movies you have watched or news article you have read on a website, for example. ‘We try to increase the diversity in these recommendations,’ Sullivan says. ‘If the recommendations are too similar to your preferences, you may end up in the well-known filter bubble. But if they are too dissimilar, you will ignore them out of hand. It is important for recommendations to be diverse, but not too diverse. If a certain user is known for only reading short news articles, you may want to recommend short news articles about slightly differing topics.’ Sullivan uses existing Natural Language Processing software, a kind of artificial intelligence, to extract linguistic features from texts. ‘In our research, we use various methods to weigh these features in order to reorder the list of recommendations. Our ultimate aim is to stimulate diverse behaviour in the user.’ Another part of their research is focussed on the best way to present the resulting list of recommendations to the user. The recommendations may, for example, be accompanied by an explanation such as: ‘If you are interested in a more in-depth understanding of this topic, you may like these articles’. ‘In our group, we are not so much interested in the commercial application of this research,’ Sullivan says. ‘Rather, we want the recommender system to represent the norms and values of both the user and the company.’

.....

‘If the recommendations are too similar to your preferences, you may end up in the well-known filter bubble.’



Peter Bosman

Optimal radiotherapy

Peter Bosman is part-time professor in Evolutionary Algorithms at the TU Delft, and senior researcher in Life Sciences and Health at the Centrum Wiskunde & Informatica (Centre for Mathematics and Informatics) in Amsterdam. Evolutionary algorithms are his specialty, both their fundamental structure and their practical application. He, for example, applies them in the

*‘We provide radiation-oncologists with a unique understanding of the competing requirements in the radiation treatment of cancer.’*

Peter Bosman



field of brachytherapy, a cancer treatment in which radiation is applied from within the patient. ‘It is often not possible to deliver the required amount of radiation to the tumour for the patient to be cured, while also limiting the amount of radiation to surrounding healthy organs in order to limit the risk of treatment related complications,’ Bosman explains. ‘Using our algorithms, the treating radiation oncologist gains a unique understanding of how these competing requirements can be balanced for each particular patient. Being the expert, they can then select the optimal radiation treatment plan for each patient, based on these insights.’ It is important for the selected treatment plan to be deliverable from a technical perspective and for it to be an accurate representation of reality. ‘We basically refrain from applying simplifications to the underlying model,’ says Bosman. In order to find a suitable solution in limited time, the evolutionary algorithm tries several possible solutions at once. Through a process mimicking natural selection, the algorithm will combine the good features of these solutions into ever better solutions. A close collaboration with radiation oncologists is important for the results to be useful in the clinic. ‘Our understanding of their underlying goals in the radiation treatment allows us to build an algorithm that calculates the desired collection of applicable treatment plans,’ Bosman says. ‘These plans are sometimes radically different from what they are used to. It prompts them to reconsider the treatment plan qualities they find most important.’

Evolutionary algorithms are a kind of artificial intelligence introduced in the ‘70s. Nowadays, it is relatively easy for anyone to apply them but, Bosman stresses, ‘they certainly are not a cure for all. Thanks to our fundamental understanding of these algorithms we are much more likely to find a good solution. For me, it is essential to combine fundamental research with its practical application.’






05

## Algorithms unravel cellular dynamics, and what it means for targeted therapy

..... Text: Merel Engelsman | Photography: Mark Prins

Modern technology allows us to monitor the activity of tens of thousands of genes in cells. But to understand how cells work, it is essential to determine how these genes interact with each other over time to carry out vital functions or adapt to changes in the environment. Novel algorithms, designed at TU Delft, uncover the fundamental building blocks of cellular processes in health and disease, and could help identify new targets for therapy along the way.






.....

Joana de Pinho Gonçalves: 'We study gene regulation and disruptions involved in cancer to identify drivers of the disease beyond genomic variation and discover new targets for treatment.'

*'As computer scientists we bring a fresh unbiased perspective to the study of biology.'*

Organs, such as the heart and lungs, are the building blocks of the human body. The organs themselves consist of many cells, working together, enabling organs to perform the tasks needed to help the body function and respond to changes in its environment. Another process of cooperation is going on at an even smaller scale, within the cells themselves. Whenever a cell has to respond to changes in its environment, certain genes will increase or decrease their activity on demand to produce the necessary

worker molecules (proteins). Very often, complex biological responses are the result of multiple groups of genes working in concert over time, with each group performing smaller well-defined tasks. 'These groups of genes can be seen as functional modules,' says Joana Gonçalves, assistant professor at TU Delft's Faculty of EEMCS. 'They may differ between cell types, as well as between healthy and diseased cells. We develop algorithms to reveal these functional modules, and their differences.'



### It's about time

'I have been intrigued by the paradigm of analysing biological systems through the eyes of modularity for some time,' says Gonçalves. 'Cells organising themselves into organs in living organisms is an example of modularity as a structural concept, in which modules are physical building blocks. In this sense, you can also think of words forming sentences in text and pixels forming objects in images. But modularity exists in function as well, such as when genes and proteins show concerted activity to perform biological tasks.' An important aspect of functional modularity is that it takes place over time. To understand the dynamics of cellular processes, the researchers analyse gene activity data measured over multiple points in time. Modern technology can do this for all genes in the cells simultaneously. The goal is to look for groups of genes showing similar activity, as an indication of functional similarity. One of the challenges is that the number of functional modules and the duration of the biological tasks they perform are unknown beforehand. 'Most available approaches, such as traditional clustering methods, group genes by looking at activity patterns over all time points,' Gonçalves says. 'This is ideal for detecting genes involved in more general biological processes, but not so much to break down the more localised coordination often found in adaptive responses. *First responder* genes, for instance, are triggered in the beginning but may not continue to coordinate as the response progresses into later stages. When we require high similarity over the entire time course, we will likely miss these effects.' Methods that can search for local patterns, however, are typically not designed for time series data. They can generate functional

modules with arbitrary 'jumps' in time. At best these are difficult to interpret and at worst they are not realistic. 'The information is in the data,' Gonçalves says, 'but it needs to be properly treated and interpreted. We needed to develop new algorithms to do this, algorithms that are time-aware and able to detect local temporal patterns.'

*'We use the properties of time to design efficient computational methods for the analysis of time course data.'*

### UDDNNNUUUD

Another challenge is that it would not be feasible to try out all possible sets of time points and search through all possible combinations of genes into modules of all possible sizes. Depending on the size of the dataset, this could take forever. 'Instead, we take advantage of the properties of time to apply useful algorithmic techniques,' Gonçalves explains. 'The fact that time has a strict chronological order enables us to establish a direction along which the search should be performed. Likewise, knowing that biological processes are expected to occur within a time window allows us to impose a reasonable restriction on finding modules with consecutive time points.' Perhaps the most important aspect of the researchers' approach is that they look at the shape of the gene activity pattern, rather than the exact values. For instance, if the activity



value of two genes increases three-fold at a certain point in time, this could be considered a match regardless of what their baseline and final activity levels were. To analyse the data, the researchers therefore determine if the change in the activity of each gene between every consecutive pair of time points is either going up (U), down (D), or not changing (N), based on the slope of the change. 'This is one among many possible discretization options that can be used,' Gonçalves says. 'Translating gene activity values into letters allows us to apply efficient string-matching techniques to find maximum-sized modules which are not fully contained by any others.' The observed string patterns have to match exactly, such as 'UDDNNUUUD,' but the discretization

criteria,' Gonçalves explains. 'We take into account the likelihood of finding each functional module by chance based on its pattern, number of genes and number of time points. For instance, we look at how often 'up,' 'down' and 'no change' occur within the pattern of a particular functional module, compared to their prevalence in the entire dataset.' These are purely data-driven criteria, independent of biological knowledge and meaning. Using biological databases, the researchers can also assess the consistency in the biological role of the genes in each functional module. 'We make use of these and other statistics to filter and rank modules, in order to select the most promising candidates for further experimental validation.'

*'Our algorithms open new avenues towards targeted therapy by uncovering groups of genes working in concert, and the regulators controlling these processes.'*

allows for some small fluctuation in the original values which makes the technique robust to noise. The researchers' most recent algorithm allows these patterns to be shifted in time to account for cases in which genes within a functional module may be activated or inhibited with a certain delay. Because of the algorithmic design strategies employed, the *LateBiclustering* algorithm can be run comfortably on a personal computer for most available time course gene expression datasets.

### Refining modules

One of the pitfalls with algorithms of a combinatorial nature is that they can generate a lot of results. 'We develop strategies to select the best functional modules, according to different statistical

### Back to biology

The researchers started off using publicly available measured data for a proof of concept. For experimental validation, they worked in close collaboration with biomedical researchers at the Netherlands Cancer Institute (NKI). 'Since the genes in functional modules show similar responses, our idea was that they might be controlled by the same regulators,' Gonçalves explains. Regulators are very interesting from a biological perspective, because they control which genes are expressed and at what rates. They are a step up in the organizational hierarchy of cellular processes, somewhat like managers supervising collections of genes. 'We identified these regulators using an approach I published during my PhD, and which I then used in combination with our



### About Joana

Joana Gonçalves, assistant professor at TU Delft's Faculty of EEMCS



*'These algorithms will allow us to discover more robust patterns'*

time-aware LateBiclustering algorithm.' The researchers used this methodology to analyse the response of human prostate cancer cells to androgens (male hormones). A number of these identified regulators were inhibited during subsequent validation experiments. Their inhibition indeed resulted in consistent changes in gene activity in the functional modules that the regulators were predicted to control. 'Together with our collaborators, we showed that our algorithms have the ability to capture meaningful biological knowledge,' Gonçalves says.

**Targeted therapy**

The analysis of gene activity over time is critical to advance our understanding of complex biological mechanisms involved in processes such as normal human growth


and development, disease susceptibility and progression, or response to treatment. Although there are still many steps to take before the proverbial 'cure for cancer' is found, the idea is that these newly developed techniques will be able to contribute to identifying new drug targets for therapy. Consider prostate cancer, for example, a hormone-related cancer that is strongly influenced by androgens. The standard treatment blocks the binding of androgens to the androgen receptor gene, which is a known master regulator of the androgen response. In many patients, however, the cancer cells develop resistance to this treatment.

It is therefore important to identify new alternative targets for treatment of this cancer. 'We have shown that our algorithms

provide a promising step in this direction,' Gonçalves says. 'By revealing how genes in prostate cancer cells organise in response to androgens, we can help understand which functional modules or biological mechanisms influence disease progression. Regulators of such mechanisms are especially appealing as drug targets, since it may be easier to target a single regulator rather than trying to directly control all the genes within an important functional module.'

One of the next steps of the researchers is to develop algorithms to identify functional modules across multiple samples. Think of cells from different tissues, different patients or different regions within the same tumour. 'These algorithms will allow us to discover more robust patterns, and for

example identify functional modules that are disrupted exclusively in cancer cells and not in healthy cells. This can then be exploited therapeutically,' Gonçalves says. 'Or we may find functional modules that show distinct patterns of activity in subgroups of patients. Based on this information, these subgroups may be suitable for alternative treatment strategies.' So far, the researchers have looked into model systems such as cancer cell lines, which are far more homogeneous than cells found in patient tumours. According to Gonçalves, 'an important challenge in drug development is to translate the knowledge we acquire from such model systems to actual patient tumours.'



**A fresh perspective**

As she had contemplated studying medicine, biology was one of the courses Gonçalves followed in high school. She chose to study computer science instead. Nearing the end of her studies, it turned out that the algorithmics professor she chose as her thesis supervisor worked on answering biological questions. Since then, bioinformatics has been her focus. Coming from such different backgrounds, it sometimes takes a bit of effort for biologists and computer scientists to understand each other. But working together allows them to go well beyond what they would be able to achieve on their own. 'Biologists have in depth knowledge

about the experimental techniques and the biological systems under study,' Gonçalves says. 'They draw hypotheses and interpret observations in light of existing knowledge. As computer scientists, with far less prior information, we offer a complementary perspective. We typically conduct unbiased data analyses, see interesting patterns emerge, and then start asking questions. It is a fresh approach to the study of biological systems, sometimes leading to surprising findings.' Gonçalves maintains active links with the NKI, and is working on joint grant proposals to expand the collaborations in a number of research directions together with groups at LUMC and Erasmus MC.



# Increased road safety with high-resolution automotive radar

06

Text: Merel Engelsman | Photography: Mark Prins

A step reduction in fatal road accidents. Whether it is by means of fully autonomous cars or the widespread introduction of advanced driver assistance systems, Professor Vaucher works on automotive radar technology for improved road safety. Equipped with millimetre wave radar, cars will be transformed into securely-connected robots with the capability of sensing the environment, thinking, and acting autonomously.

'Companies such as Waymo and Tesla, working on fully autonomous cars, get all the media attention,' says Cicero Vaucher, professor of integrated millimetre wave front-end systems at the Electronics Research Laboratory within the Faculty of EEMCS. 'But full autonomy is only the

last phase in a continuous process towards increased road safety, achieved by transferring more and more responsibility from the human driver to the car. This process started about five years ago and is now really taking off.' Road fatalities in the Netherlands peaked in the '70s at 3500 casualties per year. It was the introduction of seatbelts, followed by legislation making them mandatory, that initiated a steep decline. Thanks to the subsequent introduction of electronics systems, such as the anti-lock braking system (ABS), the airbag and car stability systems, the number of fatalities has continued to decline. 'After initially being luxury accessories, these systems are now mandatory as well,' says Vaucher.

'Within the next three years, I expect currently optional radar-based technology, such as adaptive cruise control and automatic emergency breaking, to also

become compulsory. And this is only the first step towards increased road safety and highly automated driving that automotive radar will help enable.'

*'Road fatalities in the Netherlands peaked in the '70s at 3500 casualties per year.'*

## Not yet ready for full autonomy

In this first level, the car can control steering or vehicle speed, with the human driver controlling these as well. As a society, we are already adopting the second level as well, in which control of both steering and speed is fully transferred to the car under specific circumstances, such as partially automated parking or driving assistance in a traffic jam. There even are companies offering the semi-autonomous third level of fully automated highway driving and self-parking, for which complete control is transferred to the car. 'Up to this level of automation, the human driver still carries overall responsibility for safe operation under all conditions,' Vaucher explains. 'The





driver therefore has to remain alert at all times.’ The fourth and final level, of fully autonomous driving across many driving modes – in which the driver can read, work or relax while being driven – is not yet commercially available. ‘The industry is not yet up to that level,’ says Vaucher, ‘and legislation needs to catch up. An alert and responsive driver is still legally required.’ For all levels of automation to become widely adopted, and then mandatory, it is important that the underlying technology is affordable. Lower levels of vehicle automation may require only a handful of radars pointing forward, but for the highest levels the car may need to be equipped with up to twenty radars covering all horizontal directions. ‘We have now reached the point where the hardware costs for such an advanced radar system are well below a thousand euros,’ Vaucher says. ‘For high-end segment cars this is an acceptable price level, whereas a further decrease would be welcome for cars in the entry-level segment.’


### Smaller than a passport photo

Automotive radar uses a so-called *chirping signal*, in which the frequency of the emitted radar signal increases linearly over time. This chirp is repeated tens of thousands of times each second, with a brief pause between each pulse. It allows for both the distance to and speed of objects to be determined within milliseconds, by comparing the frequency of the returned radar signals to the (then) emitted frequency. ‘Current automotive radar systems use very high frequency radio signals with a wavelength of only a few millimetres,’ Vaucher explains. ‘The three-fold reduction in wavelength compared to earlier systems, in combination with an increased bandwidth at these higher frequencies, comes with plenty of benefits. It has resulted in a much better resolution in both the measured range and velocity, while also enabling a substantial reduction in the size and energy consumption of the radar.’

### About Cicero

**Cicero Vaucher**, professor of integrated millimetre wave front-end systems at the Electronics Research Laboratory within the Faculty of EEMCS.

He shows the current state of the art in automotive radar, developed by himself and his team at NXP Semiconductors. The complete radar system consists of two tiny chips, three radar emitters and four radar receivers, all on an integrated circuit board smaller than a passport photo. ‘It can be integrated invisibly, even into small cars,’ Vaucher says. ‘This single solution can be used for all levels of vehicle automation. The difficulty with millimetre radio waves is to have the technology to reach those frequencies. But once you get there, it is a huge resource to be used.’



Inaugural speech Cicero Vaucher.



*'Adding affordable millimetre wave radar technology to all cars improves road safety and reduces road fatalities significantly.'*

**Unaffected by weather**

Radar alone is not the solution. Multiple technologies are driving the process towards increased car autonomy and road safety, with each technology having its own strengths and weaknesses. 'I believe that radar and cameras are ideal complementary sensors,' says Vaucher. 'A camera has a very high resolution for detecting objects, but it has difficulty assessing their distance and relative speed.

That is exactly what radar is ideally suited for. Radar, on the other hand, cannot read road signs, see markings on the road such as lane lines, or determine the colour of traffic lights or tail lights. Then again, radar performance is not affected by fog, heavy rain, or nightfall.' The functionality of both systems overlaps with the LiDAR system (Light Detection And Ranging) that is frequently used in car automation. It is typically placed on top of the car for an all-around field of view. Using a rotating laser, it emits millions of infrared light pulses per second. By measuring the time it takes for each pulse to return, it creates an image of its surroundings, with quite a high resolution. Some of these systems can even determine the velocity of what they see, but they are expensive, and their reliability is unproven. 'More importantly, it is not only about complementarity,' Vaucher explains. 'The more responsibility you transfer to the car, the more redundancy you want to build into your system.' Having duplicate functionality means that the car will still operate safely, even if a certain sub-system breaks down. 'We do want to increase the angular resolution of automotive radar to the point where it could replace LiDAR,' Vaucher continues, 'but it would then be up to the car manufacturer to decide on the required level of redundancy. Perhaps it is best to have both LiDAR and radar.'

.....  
Having duplicate functionality means that the car will still operate safely, even if a certain sub-system breaks down.



**\*MIMO**

By using multiple transmitters at the same time, we get the three fold increase. This is then called multiple-input multiple-output technique (MIMO). The multiple-input refers to several receivers, whereas the multiple-output refers to several transmitters

.....  
Using the passport-photo-sized chip 'as is,' an angular resolution of 36 degrees can be achieved, basically subdividing human forward vision into only five pixels. 'It will be able to detect an oncoming truck, but not the motorcycle that may be overtaking it.'




**The CRUISE project**

Improving angular resolution is one of the challenges addressed in the CRUISE project, a research collaboration between TU Delft and NXP Semiconductors. Vaucher is employed at both. He explains that 'the project addresses challenges that we know are coming. We can, for example, already anticipate the need for radar systems where you make measurements with several transmitters at the same time.' Using the passport-photo-sized chip 'as is,' an angular resolution of 36 degrees can be achieved, basically subdividing human forward vision into only five pixels. 'It will be able to detect an oncoming truck, but not the motorcycle that may be overtaking it,' Vaucher says. 'Using the innovative antenna layout in combination with smart digital signal processing, however, a three-fold increase in image resolution can be achieved by a process called (MIMO\*) beamforming.' For the overtaking motorcycle to be detected the angular resolution of the radar needs to be improved another ten times, down to a single degree. 'We know how to get there, and even beyond, by assembling four of these radar chips into a single printed circuit board. The difficulty is to get everything working simultaneously as the transmission channels will interfere with each other.'

**Tackling interference**

Automotive radar sensors share a limited frequency bandwidth. Interference will therefore increase even further as radars are integrated into more and more cars. This may result in objects with a low radar cross-section to be lost during tracking or not to be detected at all, leading to dangerous traffic situations. 'This may even become a blocking point for autonomous driving initiatives,' Vaucher explains. 'Within the CRUISE project we investigate novel approaches to tackle these interference issues. We are looking into separating the signals by coding them, for example.' Not even the powerful linear chirping signal itself is sacred in the quest to tackle interference and to realise high-resolution affordable radar. 'We work in close collaboration with experts in radar signal processing at TU Delft, such as Professor Yarovoy from the group of Microwave Sensing, Signals and Systems. Our aim is to determine the next chirp wave-form, one that would be less susceptible to other radars or self-interference. At the Electronics Research Laboratory, we have the knowledge and means to then develop the circuits that will be able to generate and decipher these signals. Automotive radar is a very exciting multidisciplinary problem and it is really important to close the loop of circuit design and signal processing.'





*'In autonomous vehicles, complementary sensor redundancy is crucial. Duplicate functionality of radar and lidar sensors means the car can still operate safely, even if a certain sub-system breaks down.'*


### Sensor fusion for optimal use of radar images

There are even more disciplines needed, as the processed signals also need to be interpreted. There is a lot of information coming from the radar; the distance towards many objects, their relative speed and at what angle they are located with respect to the car, but it doesn't tell you what each object is. The artificial intelligence used to interpret the images produced by on-car cameras will very likely also be needed for processing radar images. 'More importantly,' Vaucher says, 'different sensors need to be joined for optimal decision making. A camera might classify an object to be a pedestrian, but the radar could perhaps recognize the variety in relative speed patterns of legs and torso to be typical for a cyclist, or even for something else.'

### A long road towards fully autonomous cars

Despite rapid developments in sensor technology and artificial intelligence, it still is a long road towards fully autonomous cars. They are already being extensively tested on the road and even within cities in the United States but, according to Professor Vaucher, it may still take a few years for them to become commercially available. 'The infrastructure in these cities, consisting of broad roads and having only few pedestrians and virtually no cyclists, is almost ideally suited for initial technological development and later for commercial introduction,' he says. 'But it is a much bigger challenge for autonomous cars to navigate Delft city centre, where you have to deal with ten pedestrians and fifteen cyclists simultaneously.' Not to mention a plethora of small bridges, frequent changes in traffic situations and the erratic non-law-abiding behaviour of said pedestrians and cyclists. Asked if he will be one of the first owners of a fully autonomous car, Professor Vaucher responds that 'it is fascinating technology and it will represent a major step forward in road safety. I am looking forward to it.'



Despite rapid developments in sensor technology and artificial intelligence, it still is a long road towards fully autonomous cars.



### Millimetre wave radio signals

Millimetre waves are radio waves and, just like visible light, part of the electromagnetic spectrum. As all electromagnetic waves travel at the constant speed of light (of roughly 300,000 kilometres per second), it is the frequency of these waves that determines their wavelength. Radio waves with a frequency between 30 GHz to 300 GHz (1 GHz = 1 gigahertz = one billion oscillations per second), have a wavelength of 10 millimetres (1 centimetre) and 1 millimetre, respectively. Visible light has a wavelength that is about ten thousand times smaller.





07

## You only know the best method of defence when you understand more about your opponent

..... Text: Robert Visscher | Photography: Mark Prins

Christian Doerr conducts research into cyber threat intelligence and also devises ways of preventing new attacks.



**Your research field is digital security and you and your team explore how cyber criminals operate. Why is this necessary?**

‘ICT plays a huge role in our society. It ensures that we have water to drink, that there is food in the supermarkets and that electricity is available from the wall socket. All of this runs faultlessly because almost everything is digitised. But we have become so dependent on ICT that it’s also making us vulnerable. Imagine someone wants to do harm and attack an electricity plant or our water supply. We would be facing a major problem.’

**Why are these ICT systems so vulnerable?**

Because many of the innovations we are now using were originally designed for another purpose. The internet is the best example of this. It was devised as a network between computers that could trust each other. This means that security was not built in at the design stage. No one predicted that we would use it to connect together so many computers, in a situation in which not every user can be trusted. Besides, it’s always much more difficult to integrate security retrospectively. It should have been thought of much earlier in the design process. Unfortunately, that’s often far from being the case.

**How can you defend yourself against cyber attacks?**

First of all, you need to know what an attacker is doing. Imagine someone breaks in to your neighbour’s house. You don’t want it also happening to you. It’s likely that you’ll head to the DIY store and buy an extra lock. But by doing so, you’re assuming that the burglars will come through the front door. They may however throw a brick through the window or the criminal could have a passkey to open all locks. You only know the best method of defence when you understand more about your opponent. Because if your neighbour had a brick thrown through the window, you may be better advised to invest in some bars for the downstairs windows.

**About Christian**

**Christian Doerr** is a professor in Cyber Security and Enterprise Security at the Hasso Plattner Institute. Before he made the move to Potsdam (Germany), he worked as an assistant professor in the Cyber Threat Intelligence Lab in the Intelligent Systems department of EEMCS. Because he is still a guest lecturer at TU Delft, you see him walking around the university campus on a regular basis. About his field of research: ‘I work in the broad area of network security and critical infrastructure protection. My research focus is designing resilient network systems, localizing and estimating current threats through real-time situational awareness in networks as well as conducting threat intelligence on adversaries. Prior to joining, I was at the University of Colorado, USA, where I received my PhD in Computer Science and Cognitive Science. While most of my work focuses on technology, I also integrate socio-technical aspects of cyber security with this background in my research.’



.....

**Christian Doerr:** ‘I was mad about computers as a young boy. I found it amazing to see how a computer sends information and it suddenly appears somewhere else via a cable. It’s almost like magic!’







.....  
'We are now working on a tool that not only monitors if something is going wrong, but also assesses how serious it is.'

### Is that how it also works in your field of cyber threat intelligence?

Understanding how criminals operate is an important part of my work. What skills do they have and how do they apply them? Have you been targeted by a criminal sending out phishing e-mails in the hope that someone clicks on a fraudulent link? If so, you're a random victim. But if you have a Picasso painting on your wall at home and a burglar tries to steal it, you're dealing with a different approach. That person will do everything possible to break in and you need to protect yourself accordingly. It's a different threat profile. It also works like this with cyber threats, for example if a company has important corporate secrets. We investigate how the attackers operate and how you protect against that. If, for example, you discover that someone always uses a crowbar to break in but never throws a brick through the window, you will take that into account in the protection. Our work is similar to that.

*'Our algorithm may notice an attack on a bank that has already been tried on three other banks.'*

### What do you learn from analysing criminals?

We are working with a major telecom provider on a project. The internet is all made up of networks connected to each other. Providers ensure that these connections happen and an important part of that involves the so-called *BGP*. This Border Gateway Protocol ensures that different networks are connected to each other by exchanging information about accessibility. It is the glue that holds together all the networks, so to speak. Criminals can attack the BGP in a way that causes internet traffic to be redirected to them. Currently, we are already monitoring whether something strange is happening, but we do not yet know what exactly is going on.

### Of course, you will want to find that out. How do you approach that in practice?

We are now working on a tool that not only monitors if something is going wrong, but also assesses how serious it is. For example, our algorithm may notice an attack on a bank that has already been tried on three other banks and it's targeting accounts in Western Europe.

### You recently conducted research into how hackers operate by looking at attacks on TU Delft. What did you find?

Our discoveries included an attack by 27,000 computers all targeting TU Delft. These computers come from more than a hundred countries. They transmit information intermittently. There may be just a few hours in-between, and then several days. It just looks like random noise, but when you put it together, it's a coordinated attack. Noticing this enables us to defend ourselves more effectively. You can't stop it, but it can be blocked.

### Obviously, in the future, you will also want to stop this kind of attack. How do you do that?

One of my students is researching this. It is all about whether we can predict what an attacker will do. If he first tries this door to see if he can enter, we know that he'll move on to another door after that. You therefore install a firewall. Even better is when we put a so-called *honeypot* in that place, enabling us to capture the attacker in an environment we control. This means it's not damaging to our system and we can see what the person is doing in the honeypot. The problem is that advanced attackers, who worry us the most, have techniques that tell them they are in a honeypot. They then share that information with others. We are currently doing a study to test how they discover this. Interestingly, attackers are often better at sharing information with each other than we defenders are.

### Why is it so important for companies to share information with each other?

It enables them to warn one another. For example, attackers may be specifically targeting several banks or the transport sector. Research shows that criminals often attack many companies in a specific sector simultaneously. They also often try out attacks in Germany and France first before moving onto smaller countries, such as the Netherlands and Belgium. If German companies have already shared information about a phishing e-mail or hack with Dutch companies, the chance of them falling for it is much lower.

*'Our discoveries included an attack by 27,000 computers all targeting TU Delft. These computers come from more than a hundred countries.'*

### Are companies actually willing to share information with each other?

They can sometimes be concerned that competitors will see who they are doing business with or what they're planning for the future. This deters them from showing other companies in the same sector what type of e-mails they are receiving from hackers, for example if they're pretending to be customers. There is a solution for that. There are already ways of securing e-mail while still informing others that it may involve an attempted hack.


### How do you encourage a corporate culture in which employees share information?

By offering extra training to make them aware of threats and security. You need to reward an employee who reports a phishing e-mail. The IT helpdesk then sends round an e-mail praising the employee and warning the others. You should arrange free cake for everyone in the staff restaurant. That makes it a talking point and rewards good behaviour. People are always the weakest link. We shouldn't blame them, but enlist their help instead.

### Why are you so fascinated by your specialist field?

I was mad about computers as a young boy. I found it amazing to see how a computer sends information and it suddenly appears somewhere else via a cable. It's almost like magic! When I started studying, I became more interested in security and how attackers operate. How do you send data securely? As we have become ever more dependent on computers, this field of work, still in its infancy, has only become even more important. I hope that I can do my bit to combating the cyber threat.





08

## Protecting the electricity grid of the future

..... Text: Merijn van Nuland | Photography: Mark Prins

Power failures occur in all shapes and sizes. At TU Delft, scientists conduct research on these kinds of incidents and search for solutions how to make the electricity grid more robust in the future. One of them, Associate Professor Marjan Popov, has set up the Power System Protection Centre at TU Delft to facilitate a smooth transition to renewable energy. Using experiences from recent black-outs, he explains what can go wrong.

*'Transport chaos across England and Wales after major power cuts', 'Huge power cut paralyses Britain', and 'Urgent investigation at National Grid after worst power cuts in decade' – the headlines in the Guardian, Daily Mail and Sunday Times did not mince their words when large parts of the UK were hit by a major power failure in August 2019. Millions of households were without electricity, no trains were running, and the roads became dangerous because the traffic lights failed. A generator tripped off because of a fault, and as a result, less electricity was being supplied than used.*

It is just one example of what can happen when the electricity grid is affected by severe disturbances, according to Marjan Popov, associate professor in Electrical Power Systems at TU Delft. He and his team are working on finding solutions for rapid detection of failures and providing suitable protection against power outages. That is, as yet, far from simple, because a power outage can have various causes, from a faulty generator to a lightning strike and from broken cables to an imbalance in power production. In recent months, Popov and his team completed two projects that can contribute to a more stable and robust electricity grid. The main aim of both projects was the early detection and elimination of short circuits in the network, in order to prevent a domino effect.

*'A power outage can have various causes, from a faulty generator to a lightning strike and from broken cables to an imbalance in power production.'*



*'With a reliability level of 99.9948%, the Netherlands secures power supply almost at all times.'*

### Ridiculous mistake

'In my projects – called MIGRATE and URSES – I am looking for a way to monitor an electricity grid,' says Popov. 'We used big data research to identify the weaknesses in the system, so that we could isolate these components from the rest of the grid. An electricity grid can be compared to the human body: because everything is connected, and components which are linked together can easily affect each other.'

The fact that this can be dangerous was shown in Italy in 2003. Some 17 years later, Popov still speaks with amazement about the *ridiculous mistake* that resulted in one of the greatest ever power failures in European history. At the time, Italy imported power from France and Switzerland. But when the Italians decided to import more power from Switzerland, this proved to be too much for the Swiss power cables. It was not long before the first connections failed. Italy was without power for 24 hours, costing the country billions of euros. Only the inhabitants of Sardinia were unaffected. A domino effect could have also left other European countries in darkness if Italy had not been quickly disconnected from the European power network in order to prevent the spread of the black out.

### Renewable energy puts the cat amongst the pigeons

With a reliability level of 99.9948%, the Netherlands secures power supply almost at all times. But, according to Popov, it is not inconceivable that the electricity supply in the Netherlands could become less reliable in the future. There is one key reason for this: the high utilization of renewable energy sources. Wind and solar power in particular are supposed to cause a sharp drop in CO<sub>2</sub> emissions. However, totally revamping the structure of an electricity grid from

a handful of coal-powered plants to thousands of wind parks and solar panels will make the grid susceptible to disruption.

It is in the nature of the new phenomena that makes such a difference. In a 'green' electricity grid, fault currents are usually with lower amplitude and contain higher harmonics, and as such they may have a negative effect on the efficiency of the grid. In addition, renewable sources have little to no inertia, which means that future electricity grids will be less stable than coal-powered plants. All these disturbances are very difficult to detect by applying classical protection.

'In the past six years, my team and I have made a lot of progress in this complex research field,' says Popov. 'We have developed an efficient new algorithm and protection platform in the MIGRATE project, which can detect power outages and immediately isolate affected parts of the system. In the NWO-URSES project, we have created a new platform for collecting data in real time, which allows us to measure voltages and currents; as a result, we can swiftly detect the 'sick' part of the grid and isolate it.'

### Power System Protection Centre

'The electricity grid will become more complex after the transition to renewable energy, much more complex than most of us would imagine,' says Popov. 'As we will be taking the large power plants out of operation, we will have to think about new ways of preventing power cuts and electrical resonance. In other words, we must once again make our system immune to disturbances. Therefore, we need to come up with reliable protection solutions, and we need to do it soon.'



*'As it turns out, scientists are still unable to cope enough with lightning strikes.'*

Marjan Popov: 'The electricity grid will become more complex after the transition to renewable energy, much more complex than most of us would imagine.'

This is the reason that Popov set up the Power System Protection Centre, a TU Delft collaboration with the Dutch grid operators TenneT, Alliander and Stedin, and the companies General Electric and Siemens. TU Delft already has everything that is needed for leading research in this field. The existing Real-Time Digital Simulator and the future Electrical Sustainable Power Lab put TU Delft in the perfect position to study the challenges of the future in all their facets, and Popov will undoubtedly continue to play an important role.

### Lightning remains unpredictable

So, have Popov and his team now tackled all types of power failures? No, not yet. As it turns out, scientists are still unable to cope enough with lightning strikes. In its search for earth, lightning usually strikes tall objects, such as electricity towers. Of course, these are well protected by lightning conductors and other mechanisms, but lightning still regularly strikes transformers as well. For example, in June 2012 it happened in Nieuwegein when a lightning hit a transformer. As many as 70% of residents – a total of 17,000 households – were without electricity for almost 24 hours. This not only caused annoyance, but also financial losses, resulting from shortage of electricity. This is not the only power outage to have been caused by lightning. 'There is no way to predict the impact of lightning,' says Popov. 'In certain situations, the protection measures fail too. To optimise this protection, we need research into the interaction between lightning and the earthing system. As part of the SINTEF project, I am developing a unique interface to demonstrate how to design reliable protection.'





09

# Computer as a sheepdog

..... Text: Peter Baeten | Photography: Peter Baeten

Hybrid intelligence stimulates both human intelligence and the development of the artificial variant. Catholijn Jonker is a strong advocate of this new vision. In other words: the computer as a sheepdog for humans.



The question of what artificial intelligence is and what problems it should solve has shifted with the passage of time. In the past, people were far better and faster at calculating than any machine. Since then, the problem of calculation has been totally solved by computers. But, quite aside from this, the mainstream approach in artificial intelligence (AI) has always focused on autonomous systems that can replace people and human activities. 'However, I see much more potential in intelligent systems that collaborate with humans, can adapt to changing circumstances and explain their actions. The idea is that you make much greater progress by combining artificial intelligence with human intelligence. This enables both of them, human and computer, to grow. Compare it to a sheepdog that accentuates the sheep-herding capacities of a shepherd. We call that hybrid intelligence', explains Catholijn Jonker, professor of Interactive Intelligence at the Faculty of Electrical Engineering, Mathematics and Computer Science (EEMCS).

### Gravitation programme

This radical basic idea was last year awarded funding by the so-called Gravitation programme, financed by the Ministry of Education, Culture and Science (OCW). This funding will enable leading researchers to spend ten years working on fundamental research and collaboration. The research proposal *Hybrid Intelligence: Augmenting human intellect* was awarded a total of €19 million. Together with colleagues at Amsterdam's Vrije Universiteit (the main applicant), the University of Amsterdam, and the universities of Leiden, Utrecht and Groningen, Jonker and her Delft-based Interactive Intelligence department will set to work developing theory and methods for intelligent systems that collaborate with humans. The aim is for these systems to enhance our strengths and compensate for our weaknesses.

Jonker: 'In the Faculty of EEMCS, our main focus is on the dialogue between humans and machines, in which they each help each other by shining their own light on the shared challenge. This means having the capability to explain your thoughts to the other, and the social intelligence and ability to learn from that dialogue.'




*'I think that the rise of artificial intelligence will certainly have painful consequences in many individual cases and we need to be vigilant about that.'*

In other words, it is essentially about building the bond between the computer and people, especially for the longer term. 'Currently, for example, coaching systems work for a few weeks at most. But the interaction in a hybrid intelligence system needs to be able to adapt over a long period; and it needs to remain engaging and enthralling. This means that the development of hybrid intelligence will very much depend on knowledge from other disciplines, such as psychology.'

This is why it is particularly useful that Jonker not only works at TU Delft (since 2006) but also at Leiden University. Since 2017, she has also been a part-time professor at the Leiden Institute for Advanced Computer Science (LIACS). 'This enables me to see connections with other disciplines, in a way that's not possible at TU Delft. Within hybrid intelligence, these external connections are absolutely essential.'


### Computer as co-author

Potential areas of application for hybrid intelligence include healthcare, education and science. 'A good future example could be in my work as a scientist. You could imagine a scientist having a hybrid intelligent agent as a research assistant. This assistant would start by searching for what else has been published about a particular research subject, or checking if your specific hypothesis has already been tested. The next step would involve this kind of system in thinking about your research strategy or even brainstorming with you about the subject. The ultimate consequence of this would be that the agent would need to be credited as co-author of the scientific paper. But that's still quite some way off, by the way.' 'I would be so bold as to say that this new approach makes us quite unique in the world of artificial intelligence. Although other research programmes have touched on this vision, they haven't done it in the same way as we're doing. This also makes it difficult. Perhaps you could compare us to Baron van Münchhausen: we'll need to pull ourselves out of the swamp by the hair.'



.....  
Social robot used for entertainment or to stimulate physical activities





.....

You could imagine a scientist having a hybrid intelligent agent as a research assistant. This assistant would start by searching for what else has been published about a particular research subject, or checking if your specific hypothesis has already been tested.

## Ethics

If all goes well, the Gravitation programme will not only deliver new scientific insights and applications, but will also play a crucial role in the debate about artificial intelligence and policy relating to ethics around AI. This subject of ethics touches on an essential point for Jonker: 'For example, I'm one of the founders of the Delft Design for Values institute and I'm also involved in that other programme that received Gravitation funding: Ethics of socially disruptive technologies (SDTs).' New technological advances, such as robotics, synthetic biology, nanomedicine, molecular biology and neurotechnology, but also artificial intelligence, all have potential to bring about major changes to day-to-day life, culturally, socially and economically. But they also raise complex moral issues that call for ethical reflection.

'It basically comes down to not doing everything just because you can, which is an easy trap to fall into, especially in technology. If you don't think carefully all the time, you eventually have no idea what it is you're doing. And once you've introduced a new technology, it's very difficult to remove it from the system. Just think about a problem like asbestos and how long we've been dealing with it.'

## Pocket Negotiator


One area affected by an ethical issue was the development of the so-called pocket negotiator: an AI system to support human negotiators. 'I could have opted to equip this negotiator with the ability to recommend lying a little; after all, that could have been useful strategically, but I chose not to incorporate that option.'

The pocket negotiator is a good example of an AI development that has enjoyed years of success. Jonker was awarded a Vici grant in 2007 for the project. Since then, a spin-off company (WeGain) has developed around the technology and various customers are showing an interest. 'It works in two ways. On the one hand, the pocket negotiator helps to think about your preferences for potential outcomes in the negotiation. On the other hand, it's now becoming possible to actually ask the pocket negotiator for advice during the negotiation itself.'

## Emotions

Ethical considerations are always an issue in artificial intelligence, even in the very long term. 'For example, I sometimes think about the following dilemma: what happens if we keep making everything easier for ourselves through the development of AI? Would it make us increasingly lazier and would that ultimately become part of our DNA?'

In any case, the subject of artificial intelligence evokes all kinds of emotions, especially among the general public. One recurring fear is that artificial intelligence could make enormous numbers of jobs surplus to requirements, resulting in very high levels of unemployment. What does Jonker think about that, as a scientist? 'I think you need to draw a distinction. Are you thinking about individual cases or the fate of humanity in general? For humanity as a whole, I actually think there's nothing new under the sun, apart perhaps from the pace at which changes are now happening. Jobs will certainly disappear, but others will replace them. That's always what happened throughout history. I think that the rise of artificial intelligence will certainly have painful consequences in many individual cases and we need to be vigilant about that. What will the people who lose their jobs be able to do? AI might actually be able to help in answering that question.'




.....

What is intelligence? A dog that can smell better and thus serves us in hunting or in disaster areas, complements human intelligence in a beautiful – symbiotic – way. Can we look at computers like that? That's one of Professor Catholijn Jonker's areas of expertise. In a conversation with Jim Stolze, she talks about her research and how TU Delft pays attention to the ethical aspects of datascience through *Design for Values*.

*'What happens if we keep making everything easier for ourselves through the development of AI? Would it make us increasingly lazier and would that ultimately become part of our DNA?'*





# Intermittent computing to replace trillions of batteries

10

..... Text: Merel Engelsman | Photography: Mark Prins and C. Jason Brown

Few things change as fast as the future of computing. Most experts and large chip manufacturers agree that the world's ever-increasing appetite for computing will culminate in an omnipresent Internet of Things – trillions of devices taking sensor measurements, performing calculations and communicating with each other. Intermittent computing will not change that, but it may transform how these devices are powered, changing their operation at a fundamental level.

## Warehouses full of dead batteries

*'For a sustainable internet of things, batteries must be left behind.'*

'Batteries are the biggest threat to a sustainable Internet of Things, or to computing in general,' says Przemysław Pawełczak, assistant professor in the Embedded and Networked Systems Group at TU Delft. 'The one big advantage of a battery is that it provides continuous power, allowing for uninterrupted operation of a device. But a battery is hazardous, bulky and it has a chemical impact on the environment. Without regular maintenance or replacement, it also limits the useful lifetime of any wireless device to about a year, perhaps two years.' His words of concern are echoed by Josiah Hester, assistant professor in the Computer Science department at Northwestern University, in Evanston, Illinois (USA). 'A trillion battery-powered devices averages to 130 batteries for each person on earth. Do we really want to spend a lot of time replacing them and building these giant warehouses full of dead batteries? Rather than try and change people's behaviour, we may be able to provide some technological solution to mitigate or delay the environmental impact of computing.'

Obviously, your internet-connected alarm clock or fridge can be plugged into a wall outlet. But when you bury thousands of small sensors into concrete to monitor a skyscraper's structural integrity over its lifetime, or when you envision the future of preventive medicine to be powered by deep-tissue micro-implants, these devices need to operate, maintenance-free, for decades. According to Pawełczak 'it is unrealistic to expect improvements in battery technology to solve this. To allow near-permanent sensing at low cost and at a reduced ecological impact, we have to rethink how we design these systems. We have to let go of the concept of continuous operation. Batteries must be left behind.'



### Remember, remember, remem...


Whereas classical programming and hardware architectures for wireless devices assume a continuous power source to be present, a batteryless sensor harvests its energy from the environment – by using a solar cell, scavenging energy from radio waves or by converting thermal or kinetic energy into electricity. ‘These are virtually unlimited energy sources,’ Pawełczak says, ‘but next to providing only weak power they are unpredictable and can only power a device intermittently.’ Removing the battery therefore only works in situations where you are lax about the exact time something is supposed to happen, such as a temperature measurement or computing a moving average. ‘It also introduces a whole spectrum of new problems, threatening proper operation of these devices,’ Hester explains. ‘Buffered energy may be depleted up to a hundred times per second. Such power failures require a reboot, resetting computational progress and deleting any data stored in volatile memory.’ A new approach to computing was needed to ensure proper computational operation under these circumstances. It has led to the paradigm of (batteryless) intermittent computing.

*‘Basically, grandma herself should be able to build an Internet of Things device to monitor her pills.’*

### Intermittent computing

In intermittent computing, sensing, computation and communication tasks proceed when enough energy has been harvested to turn on the processor. Program execution will continue until the stored energy is exhausted again and the device abruptly fails. ‘The paradigm of intermittent computing is to frequently

save the program state into non-volatile memory,’ Hester explains. ‘But the more often you do that, the less time you have for local computations and the less memory you have available to store local sensor data. Writing to memory also drains the device’s stored energy.’ Most contributions in the field of batteryless intermittent computing are about how to implement this strategy on existing computer hardware. ‘It is about how to write the code in such a way that it will sustain computation regardless of frequent power interruptions, while also introducing the minimum amount of overhead,’ Pawełczak says. Having met at multiple conferences, Pawełczak and Hester discovered their shared interest in proving the feasibility of a sustainable Internet of Things; energy-neutral, low maintenance and with a low ecological impact. They decided to join forces rather than compete. Together, and with their group members, they have now addressed two major challenges in intermittent computing – how to reliably keep track of time and how to ease the burden of programming these devices.



The researchers integrated their timekeeping solution into a custom-build hardware and software platform, called *Botoks*.



From left to right: Vito Kortbeek, Jasper de Winkel, Przemysław Pawełczak and Carlo Delle Donne.

### About Przemysław

**Przemysław Pawełczak** was born in Poland. In 2009 he obtained his PhD at TU Delft on ‘Opportunistic Spectrum Access’ and subsequently joined research groups in Los Angeles and Berlin. About five years ago he changed from telecommunications to the field of batteryless intermittent computing. Now, as an assistant professor in the Embedded and Networked Systems Group at TU Delft, he mainly works on the software-side of intermittent computing, sometimes building his own hardware from existing components. He likes to do long-distance runs, intermittently of course.

### It’s about time

Timeouts and timestamps are the backbone of any computing application, in particular those designed for sensing and communication. Therefore, reliable timekeeping is one of the core challenges in intermittently operating batteryless sensors. ‘Knowing how much time has elapsed also allows a device to decide on the validity of any stored data,’ Hester says. ‘A device may experience a prolonged power failure right after taking a measurement. When it comes back online, should it transmit these stored data or first take a new and more representative measurement?’ Existing solutions for tracking time, such as ultra-low-power real-time clocks, have start-up times of up to seconds and require a substantial energy boost after each power outage. They are therefore unfit for intermittent computing. The researchers instead focussed on the use of so-called *remanence timekeepers* (see image on the previous page).

‘The major advantages of a remanence timekeeper are that it can be miniaturised and that its components cost less than a penny,’ Pawełczak says. ‘Combining various remanence timekeepers, we were able to design a solution that allows batteryless timekeeping with millisecond resolution for power outages lasting up to seconds. It also has a very low start-up time of only milliseconds and a low energy demand.’ The researchers integrated their timekeeping solution into a custom-build hardware and software platform, called Botoks. Aside of the new timekeeper, it consists of a sensor, solar panel, micro controller, radio platform and an antenna. ‘We don’t want to just do a simulation and write a scientific paper,’ Hester says. ‘We have to show that it actually works. We therefore benchmarked Botoks for computational speed, energy efficiency and memory overhead under various energy harvesting conditions.’



## Just an extra piece of code

At the moment, batteryless intermittent computing is still a niche topic within the field of computer science. 'There may be about ten groups involved worldwide, but these are all based at preeminent universities,' Pawelczak says. The topic is, however, rapidly gaining traction. For the first time ever, one of the major computer systems conferences (ASPLOS) has featured a session dedicated

to intermittent computing. 'The research community is starting to recognize this as a very interesting and solvable problem, that promises to have a huge impact,' Hester says. 'We see that in the increase in scientific publications on this topic. With the community convinced, funding agencies will follow.' For a sustainable Internet of Things to become a reality, however, batteryless intermittent computing has to become much more widely adopted by both consumers and businesses. 'It is about perception and people using these devices,' Hester emphasizes. 'Basically, grandma herself should be able to build an Internet of Things device to monitor her pills.' The researchers took a large step in that direction by developing and releasing TICS (Time-sensitive Intermittent Computing System). 'There are millions of existing computer programs out there, designed for continuously-powered systems,' Pawelczak explains. 'To run these programs successfully on intermittent devices, under frequent power failures, would normally require massive re-engineering. TICS allows these old legacy codes to run as is, handling such important aspects as memory management, the passing

of time and data-validity.' TICS also brings batteryless computing right up to the doorstep of computer hobbyists. 'They can continue to use their classical programming skills, including all advanced features of the programming language,' Pawelczak continues. 'They just have to rely on an extra piece of code that they have to upload to the system.'

## Unlimited applications

Batteryless devices open up whole new application domains, way beyond an Apple watch that you never need to plug in, should that have been the first thing to come to your mind. Think of wildlife tracking, massive-scale agriculture via augmented insects and satellites the size of postage stamps that collect atmospheric data and orbit for years. So, not even the sky is the limit. And the impact of batteryless computing may go well beyond tiny devices. 'From a scientific point of view, we design things that are the ultimate of resilience,' Hester says. 'What is harder than making a computer work while experiencing a hundred power failures per second? We therefore expect developments in batteryless computing to show up in other fields of computer science as well. After all, people in large-scale computing, in space systems and security systems, are also continuously wondering how to make their systems more robust to failures.' Although it is hard to predict its arrival, we can all look forward to a future with batteries not included. 'But for now,' Pawelczak says, 'for our work to really become generally known, we may have to prove that batteryless computing is not limited to tiny sensors for scientific use, and that it also works for larger devices such as consumer electronics. We are working on that.'



## About Josiah

**Josiah Hester** comes from an indigenous, native Hawaiian background and grew up caring about the environment. In 2017 he obtained his PhD on 'Sophisticated Batteryless Sensing'. He is currently an assistant professor in the department of Computer Science, the department of Electrical and Computer Engineering and the department of Preventive Medicine, all at Northwestern University in Evanston, Illinois, USA. He works on both the software and hardware side of intermittent computing because 'codesign allows you to make something small, lightweight and perfectly tuned to the problem you are trying to solve.' His favourite areas of application are preventive medicine and large-scale smart agriculture.

## Remanence timekeepers

A remanence timekeeper essentially consists of a capacitor and a resistor. A capacitor is a device that can store electrical charge, a resistor is a device through which an electrical current can flow. You can think of them as a bucket and a faucet connected to the bottom of the bucket. Using available electrical energy, the capacitor is charged (the bucket is filled). With more electrical charge stored, the voltage across the capacitor increases (the water level rises). When power is lost, the capacitor will discharge through the resistor (water will flow through the faucet, emptying the bucket). The change in voltage across the capacitor (change in water level) once power returns is a measure for the time elapsed between two reboots. By varying the size of the capacitor (bucket) and resistor (diameter of the faucet) it is possible to optimize the timing characteristics of a remanence timekeeper.




SCAN!



During ASPLOS, Jasper de Winkel (one of the members of the TU Delft team) gave a presentation about reliable timekeeping for intermittent computing. ASPLOS is a multi-disciplinary conference for research that spans the boundaries of hardware, computer architecture, compilers, languages, operating systems, networking, and applications. ASPLOS provides a high quality forum for scientists and engineers to present their latest research findings in these rapidly changing fields.





# 11

## Developing a digital twin for the electricity grid

..... Text: Merijn van Nuland | Photography: Mark Prins

The rapid transition to renewable energy threatens to cause major problems to the very expensive electricity grid in the Netherlands. In his quest for solutions, Professor Peter Palensky is now working on a *digital twin* to make it possible to study the grid effectively.

Imagine you have been working quietly for years in the confines of a large academic institute, and, all of a sudden, dramatic changes turn your familiar world on its head. The spotlights are pointing in your direction and suddenly all eyes are on you. What would you do?

This is exactly what happened to Peter Palensky, professor at TU Delft. His specialism – intelligent electricity grids – has suddenly become a hot topic in recent years. The rapid transition to renewable energy has raised an important question: is our existing electricity grid capable of withstanding such far-reaching changes? Palensky and his colleagues need to answer that key question to prevent the move towards increased sustainability from faltering.

### Everyone an energy producer

For around fifty years, everything remained more or less the same. Every country had several large energy plants that delivered electricity on the grid for millions of customers, in other words citizens and companies. But the emergence of renewable energy has suddenly added countless new energy

*This green revolution brings some huge problems of adaptation with it. Connecting a wind farm can completely change the dynamics of a country's electricity grid.*





## About Peter

Peter Palensky (Austria, 1972) has been professor of Intelligent Electric Power Grids at TU Delft since 2014. Previous positions included principal scientist in the Complex Energy Systems research group at the Austrian Institute of Technology. He has also held posts as assistant professor or researcher at the universities in Pretoria, Vienna and Berkeley. He is currently editor-in-chief of the IEEE Magazine on Industrial Electronics.

Palensky is certainly not the type to shy away from a challenge like that. Quite the contrary. 'As a scientist, it's actually quite a privilege to find yourself in this position,' he says. 'It's as if we are at a turning point in history and have a real chance of changing the world for the better. It's an absolutely huge responsibility, but one day we may be able to say to our children: we did it.'

producers, because every wind turbine or solar panel is actually an energy plant in itself. All of a sudden, the citizens are no longer customers, but actually the suppliers! Very variable suppliers, as well, since wind speeds and solar power tend to fluctuate.

'This green revolution brings some huge problems of adaptation with it,' says Palensky. 'For example, think of the countries that have to connect large wind farms to the existing electricity grid. Suddenly connecting a large amount of new capacity can completely change the dynamics of a country's electricity grid, even leading to a power failure in the worst case.'

This means that research is required. A great deal of research. Will the Dutch grid be capable of withstanding a new wind farm off the coast of Zandvoort? How will we keep everything stable when, in the near future, millions of solar panels have been installed by individual house owners? And how can we ensure that everyone will be able to charge their electric vehicles in the future?

## Digital twin

Palensky hopes to answer these questions with the help of a digital twin, a digital copy of the Dutch electricity grid. On a digital copy, you can experiment and research to your heart's content without the risk of breaking anything. 'It's actually a dream version of the real grid,' he says. He adds with a laugh: 'Obviously, we're not allowed to tinker with the real electricity grid, so we have to do it this way.'

Palensky and his colleagues have actually borrowed the idea from large companies. Aircraft and car manufacturers have long been using digital twins to test their planes and vehicles. Examples include simulators that use 3D images, sound and movement to replicate a real car ride or aircraft landing. Palensky's digital grid is similar, but obviously without a driver's seat and seat belts.

One small digital twin is already operational. The real-time digital simulator (RTDS), packed with software full of mathematical calculations, can simulate a quarter of the Dutch grid. In the near future, this version is set to be replaced by a digital copy of the whole

*'We have a moral obligation towards the younger generation to combat climate change. Technology is an important way of achieving that.'*

network. Before that, the existing High-Voltage Lab at TU Delft will need to be converted into an Electrical Sustainable Power Lab, the ESP Lab for short. In it, it will be possible to combine crucial elements from the electricity grid, such as high-voltage masts, wind and solar energy, energy storage and distribution networks, into one functioning whole.

Palensky's first objective will be to investigate how the existing Dutch electricity grid responds to stimuli or shocks, such as new wind farms or foreign cyber attacks. But, in the near future, he also intends to do what scientists love best: innovate. 'I'm interested in finding out if we can develop an alternative electricity grid that's more adapted to the modern age. It's very unlikely to replace the existing grid any time soon, but I think we have a lot to learn from this theoretical exercise.'

## Hard facts from TU Delft

The digital twin will provide grid operators, such as TenneT, with a lot of information on how to make and keep their networks stable. Palensky believes that the

'hard facts' from TU Delft will prove crucial because it would otherwise take much longer to resolve these kinds of complex infrastructural issues. Besides, this involves sensitive information that it is advisable not to obtain from distant countries: recent cyber attacks abroad have shown that the electricity grid is a popular arena in which to wage an electronic war.


Palensky and his colleagues are currently working on a proposal for funding from the Dutch Research Council (NWO), to enable them to complete their digital twin. 'It would be a terrible waste if we can't make progress now. TU Delft is a world leader in this field and we see on a daily basis how much Dutch grid operators can benefit from our knowledge.'

So what is Palensky's secret weapon? 'My students. They're the ones who often have the smartest and most creative ideas. I also think it's important to give them access to the best machinery possible. What's more, we have a moral obligation towards the younger generation to combat climate change. Technology is an important way of achieving that.'

The digital twin will provide grid operators, such as TenneT, with a lot of information on how to make and keep their networks stable. The 'hard facts' from TU Delft will prove crucial.







## How smart sensors can prevent epilepsy

..... Text: Koen Scheerders | Photography: Mark Prins

In Delft and Rotterdam, Wouter Serdijn and Christos Strydis are collaborating on a network of sensors and stimulators for the body. By picking up signals and sending the brain a rapid wake-up call, they hope to be able to predict and prevent epileptic fits. 'If we can close the loop, we'll have the technology ready within three years.'




Epilepsy is a collective name for disorders of the brain caused by a change in the brain cells' electrical activity. These neurons suddenly become overactive, firing their electric signals uncontrollably. This sometimes results in a 'short-circuit': seizures where a patient loses consciousness or has involuntary muscular spasms. For the more than 200,000 Dutch people with epilepsy, this can be very dangerous. But what exactly happens during an epileptic seizure? What exactly do the overactive brain cells do? And how can we influence that process? Wouter Serdijn, professor of Bioelectronics at TU Delft, is trying to answer these questions by researching and designing technology for monitoring, diagnosing and treating epilepsy.

### Get-together

'Do you know the band Rage against the Machine?' asks Wouter. 'I always compare epilepsy to their song *Killing in the name of*. Epilepsy experts say I'm mad for making the comparison, but that's what the signals remind me of. In the chorus, everyone's jumping all over the place.' That's what happens with brain activity during a seizure. All the neurons synchronize, reducing the exchange of information to zero. The result is an epileptic fit.

That kind of seizure can really affect patients' quality of life, says Wouter. 'All of those 200,000 Dutch people are prohibited from driving or using heavy machinery.' They also need to exercise constant caution in their day-to-day lives: stress or flashes of light can trigger a seizure. In addition, drugs are not always successful in preventing seizures. 'In major seizures, the brain cells simply break down. They literally become exhausted and die.' Wouter is therefore applying a technological perspective in his search for solutions that could help the body prevent these 'wayward get-togethers' between brain cells.



.....  
'In the chorus, everyone's jumping all over the place.'

### About Wouter and Christos

'It's actually totally logical for neuroscientists and engineers to join forces,' says **Wouter Serdijn**. He is professor of Bioelectronics in TU Delft's Faculty of Electrical Engineering, Mathematics and Computer Science. He has spent his academic career in Delft: from his Master's degree to his appointment as professor. 'I'm a Delft boy. In some ways that's actually quite boring. But I just think TU Delft is an amazing place.' Wouter is one of the few engineers with an innate interest in medical technology. He applies his expertise in bioelectronics to his work on such areas as pacemakers, hearing implants and neurostimulators.

'You call it boring, but you have quite a few miles on the clock,' says **Christos Strydis**. He did his Master's degree in Computer Engineering in TU Delft where he also completed a doctorate. His career in neurosciences took off in Rotterdam, first as a postdoc and now a tenured assistant professor with the Neuroscience Department of the Erasmus MC. Since then, he has worked as an assistant professor in the Neuroscience Department at Erasmus MC in Rotterdam. As an engineer there, he forms a bridge between neuroscience and technology. Together with Wouter, Christos was recently awarded a grant by the Delft Health Initiative to develop the ECLEPSys project as part of Convergence, the joint programme of TU Delft and Erasmus MC in Rotterdam.





*'In a fit, all the neurons start jumping at the same time. In major seizures especially, they literally become exhausted.'*

### Wake-up call

As professor of Bioelectronics, Wouter is working on technology that interacts with the electricity in the body, attempting to measure and influence it. But, for epilepsy, that's easier said than done: just like with the crowd jumping up and down, you have little control over it. The entire cerebral cortex – the area of the brain that receives, interprets and analyses information – is involved in the process, says Wouter. 'So, what you need is a wake-up call for the whole cerebral cortex that can restore normal communication without bringing it to a standstill.'

Together with the Erasmus MC in Rotterdam, he is conducting research on the role of the cerebellum in conveying information. 'The cerebellum regulates the body's motor movements and it branches into all parts of the brain. These connections can help us close the loop: accurately detecting where and when a seizure develops, making an intervention that can reach all parts of the brain and then measuring the effect of that intervention.'

### Rocket science with patients

Christos Strydis is one of the researchers of Erasmus MC who is bridging the gap between the clinic and the technology. A computer engineer with roots in Delft, he works as an assistant professor in the Neuroscience department. 'I sometimes feel like a translator,' says Christos. 'In Rotterdam, there's a lot of knowledge about brain function and how the different parts of the brain interconnect and the people in Delft have a lot of experience in building technological solutions.' Christos has to speak both languages. 'Some of the bioelectronic solutions we devise for neurosciences are not rocket science, but others are. If we stayed in Delft, we engineers would never be able to gain experience with patients or, we would never really understand the essential minutiae actual neuroscientists face.'

With that in mind, Wouter and Christos are setting up a platform to explore the medical issues surrounding epilepsy from an engineering perspective. They have received a grant from the Delft Health Initiative for their work. Their project ECLEPSys – Ensemble, Closed-Loop, Epilepsy-Prevention System – consists of a prototype of a sensor network made up of implantable and wearable sensors and stimulators. The first group measures signals from the brain and the second links a signal back into the body. The idea is that the signal should stimulate the brain cells in such a way that they revert back to normal.




*'In Rotterdam, there's a lot of knowledge about brain function and how the different parts of the brain interconnect and the people in Delft have a lot of experience in building technological solutions.'*

### Sandbox

ECLEPSys involves building a medical body area network (MBAN). This is a partly wireless network made up of wearable, injectable and implantable nodes, each with its own task: sensors, stimulators and a computer that coordinates the signals with each other. An MBAN is more effective than existing technology because it is capable of measuring on or next to the body and intervening in real time, tailored to each patient's particular needs. ECLEPSys can therefore make a difference in epilepsy in particular, but Christos has even bigger ideas. 'The MBAN could also be used to treat other disorders: heart arrhythmia, Parkinson's, migraine or tinnitus. We aim to demonstrate that it is possible to help patients by applying machine learning and smart algorithms to biosignals.' ECLEPSys can be compared to a sandbox, where diverse ideas and novel technologies can be tested. 'Examples could include new sensors, new forms of stimulation and ways of analysing personal data. But the project also encompasses data security, which has already become a nemesis in the modern digital world.'





*'If we engineers all stayed in Delft, we would never have experience with patients.'*

For epilepsy patients, it is hoped that ECLEPSys will bring about a new treatment method, combined with drugs. 'I expect that this will enable us to suppress or control epilepsy,' says Wouter. 'The great thing about this type of neurostimulation is that you just switch it on and it works immediately. And vice-versa: you stop if you no longer want it. That makes treatment personal and targeted. Drugs can never be so fast: they don't know where in the body they're supposed to go.'

#### Work to be done

But that is still some way off, says Wouter. 'We already have a system that can detect and suppress epileptic fits but it only works with mice at the moment. Their cerebellum is very similar to that of humans.' By measuring brainwaves and stimulating the cerebellum, Wouter and Christos were able to suppress an epileptic fit elsewhere in the brain. That demonstrates that the principle works, says Wouter. 'We were able to detect that a fit was coming within 0.4 seconds and stop it.'

By measuring brainwaves and stimulating the cerebellum, Wouter and Christos were able to suppress an epileptic fit elsewhere in the brain.


*'There is significant variability between patients. That calls for a different approach.'*

But before epileptic patients can start driving again, there's still a lot of work to be done. The energy management of the sensors and stimulators requires improvement, for example. The technique that Wouter and Christos are using has still not been cleared for use on humans. 'Our system is still not a network. We're working on other ways of getting signals into the brain. Instead of EEGs, we're working on different types of feedback. We know that the body sends signals before a fit starts: sweating, dilated pupils or an increased heartbeat. We need to learn to use that information.'

#### Translating

Before Wouter and Christos can use the system on patients, they aim to make it as efficient as possible first. The collaboration between Rotterdam and Delft is essential in this: the realism of the medical experts in Rotterdam keeps the Delft engineers and their get-things-done mentality in check. 'We would never have made such progress without our contact with Erasmus,' says Wouter. 'Engineers believe in the power of repetition. If we make something twice the same way, it also needs to work the same way. That principle does not work on patients. There is significant variability between patients, and even in the same patient at different times. That calls for a different approach.'

Besides that, it's not only about the technology, adds Christos. 'Before rolling this out to more people, you have to discuss such issues as ethics and quality-of-life. You cannot measure brain activity everywhere in the body. Do you really want to burden a patient with all kinds of sensors on his head?' Despite this, Wouter thinks that the technology behind ECLEPSys can be translated for human use within the foreseeable future. 'You cannot apply substandard work on people. But I'm still optimistic. If you can predict epileptic fits with the right signals, we can close the loop. Then we'll have the technology up-and-running within three years.'



Wouter Serdijn was a guest lecturer on the Dutch show *Universiteit van Nederland*. Scan the QR-code to hear Wouter explain his research (in Dutch).



# Modelling a flexible electricity system

13

..... Text: Dave Boomkens | Photography: Mark Prins

In most cases, electricity is transported from point A (where it is generated) to point B (where it is used). Huge quantities of electricity are sent over the high-voltage national electricity network, while the electricity that comes out of the plug socket at home is part of the low-voltage distribution network. Using numerical analysis, TU Delft PhD student Marieke Kootte is developing a model to link these two networks. 'Because,' Kootte explains, 'the energy transition, which is having a huge impact on our national grid, makes it increasingly important to also focus on the low-voltage distribution networks in more detail.'

It is hard to imagine a landscape with no power lines, but have you ever thought about the route that the electricity takes through the thick cables before eventually reaching your iPhone? Marieke Kootte didn't really think about it either, until a call was released by the Dutch Research Council (NWO) for research to give the energy transition a push in the right direction using computational intelligence (the learning ability of computers). 'This closely matched the expertise available in Professor Kees Vuik's team, which is high performance calculations in large electricity networks,' explains Kootte. 'I decided to take a closer look at the subject, and found that there is still much to be done when it comes to coupling low- and high-voltage electricity networks. That's how the whole thing started.'

## Generated energy

Marieke is now a PhD student and the NWO call forms the basis of her research. In this research, she also makes use of knowledge gained in earlier projects. 'I once took part in a project to find out how quickly drug waste enters the sewerage system. I simulated how the mix of chemicals spreads in 2D, and did the calculations on my laptop. But to do this for the whole of the groundwater system in the Netherlands, you would really need to use numerical analysis. Using computer

*'There is still much to be done when it comes to coupling low- and high-voltage electricity networks. That's how the whole thing started.'*



### About Marieke

TU Delft PhD student **Marieke Kootte** is developing a model to link two networks. 'Because,' Kootte explains, 'the energy transition, which is having a huge impact on our national grid, makes it increasingly important to also focus on the low-voltage distribution networks in more detail.'

models that you develop using numerical analysis, you can map this kind of thing really quickly, saving a lot of time and computing power. For this study, I am applying the same trick to the electricity grid in the Netherlands: a network that is made up of thousands of cables and connections. Each cable has a certain resistance, and a connection is a point at which energy is either generated or used, such as a house, or a whole neighbourhood. Once we have mapped the details, we can run a powerflow simulation to obtain a voltage profile. This gives an idea of what is happening in the grid. It provides useful information for the grid operators, especially if they want to build a new network or connect a renewable energy source to the grid.'

### Sharing is caring

The electricity network is so big that simple models (which describe a few connections) are often not enough, and can give a distorted picture of reality. So, for modelling at such a large scale, numerical analysis really is the best option. 'But even so,' says Kootte, 'there are still some complicating factors. One example: the generated energy is transported over long distances through the national grid and transformed to a lower voltage at certain points in the network for use in homes and industry. What makes things difficult, is that each of these networks has its own grid operator. These grid operators each map the distribution networks in their

own way, and they are hesitant to share this 'sensitive' information. When I started my research, this was a huge problem. Luckily, now that the energy transition is taking off, more information is being shared between the operators.'

### Flexibility

In Delft, Kootte focuses mainly on analysing the electricity network, but 70 kilometres away – at Sympower, an Amsterdam start-up – she is working on another part of her PhD research. Kootte: 'If there are fluctuations in the supply and demand of electricity in the network, gas- or coal-fired power plants are usually used to make up the difference. That is often expensive, and it's not very sustainable. If there is a shortage or surplus of electricity in the network, Sympower switches appliances that are not directly needed on or off, or higher or lower. They therefore play with the demand side of the electricity network. For example, lots of cold storages in the Netherlands are set at 7°C. If there is a shortage of electricity somewhere, then Sympower could switch off those cold storages (in participating supermarkets) for a few seconds. As a result, the temperature may increase to 8°C, for example. The demand for electricity decreases locally, and balance is restored to the network. Controlling the storages remotely in this way means that no polluting resources need to be used to restore the balance.'







### Finnish greenhouse horticulturists

A flexible electricity system that matches demand to supply could make an important contribution to the energy transition. Testing such a system does not stop at the Dutch border. For her research, Kootte ‘controlled’ the lamps in greenhouses for growers in Finland. ‘Plants need water, food, humidity and a minimum temperature, but of course also light to grow. It does not really matter at what time the plants absorb this light during the day, as long as they get a certain minimum amount of light. That means, enough light for photosynthesis to take place. The plants do however need a certain amount of light each hour to stop them from ‘falling asleep’. So what we did was to predict the daily energy price fluctuations as precisely as possible. Then, during the most expensive hours, we let the plants sleep, and turned up the lamps to maximum during the cheapest hours. No more than a simple optimisation issue.’

### To the emergency room

Another optimisation issue that Kootte has worked on relates to electric cars. What is the most efficient way to charge, let’s say, a group of 8,000 cars? ‘You can predict when each driver leaves or returns home, but you can also calculate how many cars leave and return each hour. If the electricity network is experiencing a peak load, you let less electricity out of the charging point so that the car charges more slowly. And therefore prevent overload on the grid. Of course, you need to build in a safety margin so that you don’t have a situation in which you need to drive to the emergency room at 2 o’ clock in the morning only to discover that the battery’s not even half full. Nobody wants that.’

### Back to the future

Back to Delft, where Marieke is connecting high-voltage and low-voltage distribution networks. How exactly does she do that?

‘There are two ways to do it. One way is to assume that grid operators share all their data with one another. The other is to assume that grid operators keep all their information to themselves, except when it concerns a common connection. In the case of the second scenario, you can use a simple model for the calculations.’ However, notes Kootte, at the same time more and more energy is being consumed and more is being produced locally. This increases the level of detail required, and therefore the size of the models needed. Furthermore, energy from sun and wind is inherently unpredictable. ‘If you then link the high-voltage network to the low-voltage network, the model becomes pretty complex, with more than a million connections. The best thing then is to use a GPU – a special processor that is designed to process graphical data. A GPU can process multiple tasks simultaneously, so that the analysis can be done more quickly. A GPU is used differently from a normal computer (CPU), and that is where we – mathematicians – come into the picture. What my goal is? To apply GPU processing to increase the speed of the software that grid operators use, so that they can respond more effectively and so that electricity is always available, with a simple push of a button, even in this rapidly-developing world.’

### Huge laboratory

Marieke is not the only one working on this subject at the Faculty of Electrical Engineering, Mathematics and Computer Science (EEMCS).

A huge laboratory is currently being built at TU Delft – the ESP Lab – to investigate the best way to transport electricity produced by wind turbines and solar panels to the plug sockets in our homes. One of the people involved in the ESP Lab is Peter Palensky. He is developing a digital twin for the electricity network, to answer the question how to keep everything stable when, in the near future, millions of solar panels are installed by individual homeowners. ‘Peter is specialised in intelligent electricity networks,’ says Marieke. ‘It therefore makes sense for us to meet up now and again to discuss the issue. He knows everything about the network, and we do the superfast analyses.’

### Well-known pass

It is not just in her work that Marieke tries to make sure that we use the earth in a way that ensures its continued use for future generations. One of the things that she does, in addition to her PhD, is make sustainable clothing. ‘On Tuesday evenings, I follow a seamstress course,’ she says with a grin on her face. ‘That means, making new clothes and altering existing clothes. I love doing it, taking along my own fabrics, learning the old-fashioned techniques... I am also a pretty fanatical cyclist. I take part in the Amstel Gold Race every year, and I have cycled the Liège-Bastogne-Liège race several times, the final race of the spring classics. This summer, I wanted to cycle the Stelvio pass, a well-known climb in the Italian Alps.’

A flexible electricity system that matches demand to supply could make an important contribution to the energy transition.






## New Electrical Sustainable Powerlab smooths the way for energy transition

..... Text: Robert Visscher | Photography: Mark Prins | Infographic: Dirma Janse

A new laboratory, the only one of its kind in the world, is being built in Delft. The Electrical Sustainable Powerlab will bring together under one roof scientists researching the generation, transfer, distribution and use of electricity by households and companies. The aim is to ensure a smooth transition to more sustainable energy.



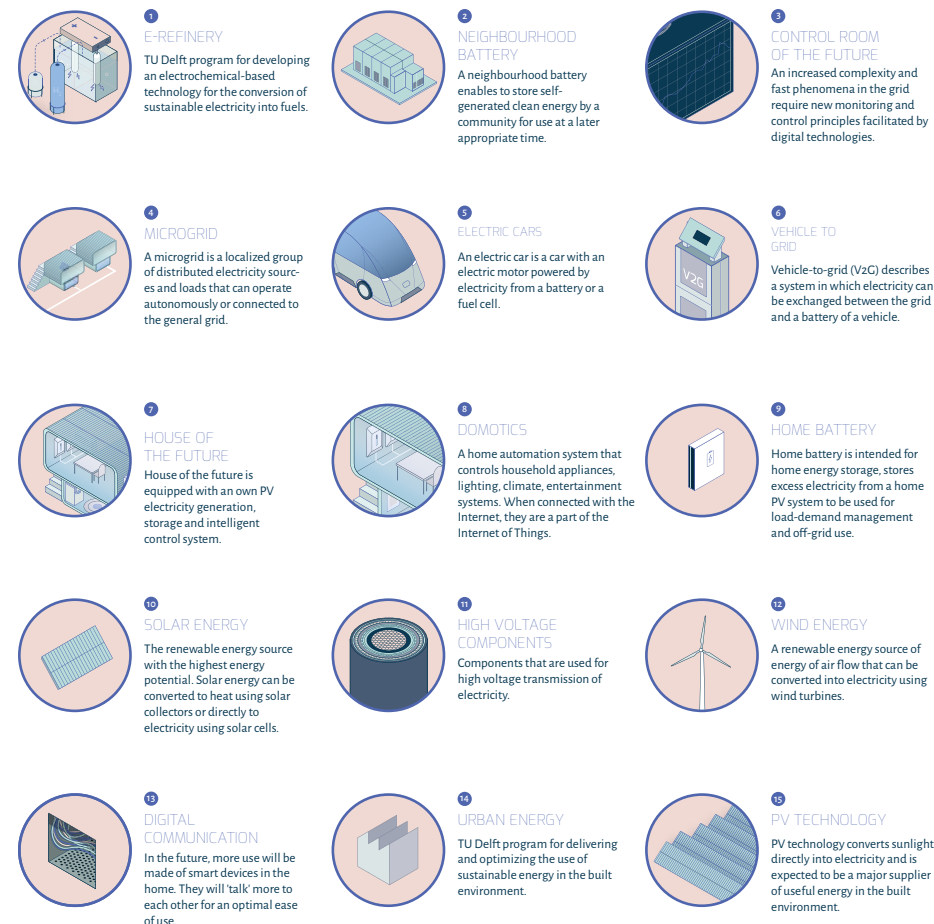
*'A power outage can have various causes, from a faulty generator to a lightning strike and from broken cables to an imbalance in power production.'*

In a large hall, exploring the most effective way of getting electricity from wind turbines and solar panels across the country to the wall socket in your home. This will be the mission of the new laboratory in the Faculty of Electrical Engineering, Mathematics and Computer Science, which will be operational in the beginning of 2021.

According to TU Delft Professor of Photovoltaic Materials and Devices Miro Zeman, there is a real need for this new Electrical Sustainable Powerlab (ESP Lab). 'The electricity system as we now know it is set to change completely,' he says. Currently, most electricity is generated centrally at coal-, nuclear- and gas-powered plants. They provide a stable supply of energy to the electricity grid.



## ACTIVITIES ESP LAB



## Digital Energy

We are already seeing increasing numbers of sustainable generators of electricity, such as wind turbines and solar panels. According to Zeman, that will only increase in the years ahead. 'People with solar panels can supply power back to the grid. This means that power is being generated not only centrally, but also at local level. It's not only controlled by big companies, but by ordinary people like you and me. This two-way traffic is an important factor that we need to take into account.'

For the electricity grid of the future, the fact that the supply from sustainable generation fluctuates also matters a lot. If it's very windy or there's a lot of sunshine, more energy will be available. 'This has made the supply of

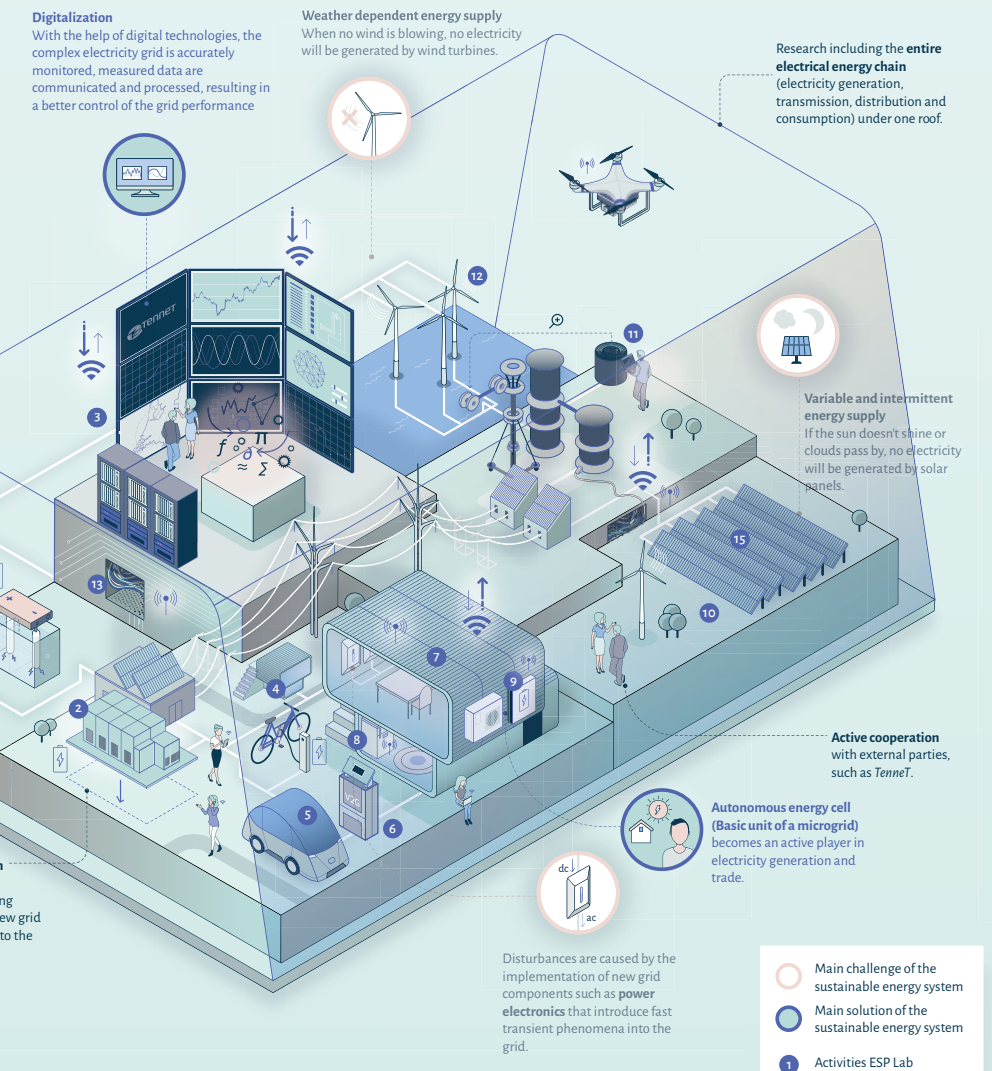
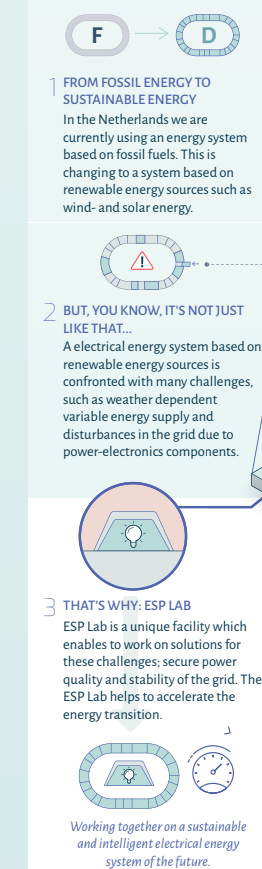
electricity more unpredictable, so we need to use more measurement and control technology. Ensuring everything runs smoothly is becoming a more complex process. That's why we're developing new methods to gain greater control of electricity generation and consumption. It involves collecting a great deal of data. We use the term digital energy to describe this new approach.'

This change will come about in the coming years. Currently, research work on the energy transition is done in separate labs in the department. For example, work on transport over long distances takes place in a High-Voltage Hall and, in another lab, scientists are testing components as converters for the transport of electricity.

## TU Delft Electrical Sustainable Power Lab

First laboratory in the world enabling research and education in system integration of new components and electricity-generating technologies into the existing electrical grid for the acceleration of energy transition.

### ENERGY TRANSITION, HOW DOES IT WORK?



## Electrical Sustainable Power (ESP) Lab

Delft University of Technology is working on one of the major societal issues of our time; the transition towards sustainable energy. With its MSc and PhD educational programmes TU Delft is preparing young engineers both nationally and internationally to become future energy leaders in realising this far-reaching change in the energy supply and market. Industry and academic experts are working together on innovations that contribute to a higher efficiency for the electricity generation from renewable energy sources, electricity transmission & distribution and system services without compromising the reliability of energy supply.

In order to tackle the challenges of energy transition, multidisciplinary research programs based on system integration are required. The Electrical Sustainable Power Lab will be a unique location where these research programs will be executed.



## About Radek

‘Over the years, I’ve worked at electricity companies and seen a lot of laboratories. The new Electrical Sustainable Powerlab will take unique position among other laboratories due to integral research approach covering the whole energy chain. That’s why I call it the laboratory of the future,’ says the current manager of the High-Voltage Hall, **Radek Heller**. He will also be the manager of the new ESP Lab.

Heller is most enthusiastic about the sharing of ideas between the different researchers. Several research groups will soon be working in the same lab. ‘At the moment, there are lots of scientists working on different pieces of the puzzle, but we’ll soon all be solving the whole puzzle together. I’m also looking forward to the mutual conversations between the researchers in the lab. People from different departments will soon be able to chat there.’

He is also keen to emphasise the important role set aside for the business community. The scientists aim to develop new solutions in the lab that will go on to be used by real companies. ‘Businesses will also need to see the benefits of the solutions we devise. We can come up with the most ingenious devices, but if they’re not used, nobody gains anything. We want to develop things that really benefit people and help society to progress. This new lab will certainly contribute to that.’

## About Armando

In the coming years, new, large wind farms will be developed off the coast of the Netherlands. They will produce direct current, but the electricity grid uses alternating current and that means losses over long distances. Interestingly, most household devices use direct current, including TVs, tablets and smartphones. ‘We’ll soon be using the new lab to explore whether it’s possible not only to

transmit but to distribute energy under direct current. The energy distribution at medium voltage direct current levels at large scale has to be done in a safe, efficient and reliable way,’ says TU Delft scientist **Armando Rodrigo Mor**. He specialises in components that work at high voltage.

In the near future, Rodrigo Mor expects to see two electric sockets in the home: one for direct and one for alternating current. But that is still some way off. The reliability of transporting high-voltage direct current over long distances still needs to be improved. ‘The new lab will prove particularly suitable for that because we’ll be able to see what works most effectively for the whole system’, he says.

The lab will be operational in the beginning of 2021.



Radek Heller, Miro Zeman and Armando Rodrigo Mor.

*‘We’ll now be bringing all of this research together in the ESP Lab.’*


## From generation to the wall socket

‘We’ll now be bringing all of this research together in the ESP Lab,’ says Zeman. ‘This will involve looking at the whole system. Imagine we’re devising a new piece of equipment, such as a transformer or converter. The new laboratory will make it possible to test how everything connected to the electricity grid will react to it. In order to check that no faults develop and to see if it works the way we think it does. In other words, we’ll be testing the whole system rather than just the performance of the individual components. This is why the Electrical Sustainable Powerlab will be essential in smoothing the way for the energy transition.’

Researchers will soon be able to investigate which cables are most effective, from generation through to the wall socket. Or to see what the effect is of converters that convert direct current from wind turbines and solar panels into the alternating current on the electricity grid. ‘Other examples could include a type of converter that enables energy from your solar panels to be stored directly in the battery of your electric car. And if there is not enough electricity in the evenings, you can use the same energy again to keep the lights on. There’s so much research to be done on everything from generation through to the user and we can do all of that in the new lab.’

*‘We’ll be testing the whole system rather than just the performance of the individual components.’*






## Understanding noise – from quantum fluctuations to climate models

..... Text: Merel Engelsman | Photography: Mark Prins

Noise adds arbitrary fluctuations to both physical and economical processes. As a consequence, there are many abrupt variations in signals and measured values that should be either constant or varying only relatively slowly. Although noise itself is unpredictable, its influence can certainly be modelled. Using the correct mathematical insights, these models can then be solved.

*'The apparent random motion of pollen in water is a good way to model noise in, for example, the realm of economics.'*







Inaugural speech.

### Noise at any scale

As professor in the Analysis group at the Delft Institute for Applied Mathematics (DIAM), Mark Veraar has a wide range of mathematical interests. Noise has been one of these interests ever since he was a PhD student. 'Noise is everywhere, at every scale of observation,' he says. 'From quantum fluctuations at the level of elementary particles, through random motion of dust particles in a gas or a liquid, to turbulence in the atmosphere and in climate models. Noise is not only a physical phenomenon, it is also a characteristic of economic processes such as a sudden rise in the financial markets or the fluctuations in the price of shares. Rounding errors in a computer are yet another kind of noise.' It is fair to say that our world is a noisy place. More importantly, noise can have a decisive influence within the models that engineers and other scientists use to make predictions about this world. The few meters rise in sea level as predicted by climate models will most likely be the consequence of (modelled) human activity. But it could, in a manner of speaking, also be due to a seemingly harmless rounding error running out of control. 'By specifically modelling noise,

*'The more we know about the smoothness of a function, the more efficient we can approximate it with numerical methods.'*



a model may more closely conform to reality. It also allows us to make statements as to the likelihood of such unwanted behaviour,' Veraar explains. 'It is for this reason that I also investigate the long-term behaviour of model solutions. Will an equilibrium be reached, and at what timescale?'

### Hopping pollen

'I use the mathematics of the so-called *Brownian motion* to model noise in problems in the realms of physics and economics,' Veraar says. The Brownian motion is the phenomenon that small particles, such as pollen, display a seemingly arbitrary pattern of motion when immersed in water. Their erratic bouncing around is a consequence of many random collisions with the even smaller water molecules. The mathematics of the Brownian motion is over a century old, but new features continue to be discovered. 'When modelled as such, noise is unpredictable and it will have a uniform intensity,' Veraar says. In a climate model, for example, this means that the impact of noise does not depend on where you look (the location) or when you look (at what modelled time).



### About Mark

Mark Veraar is professor in the Analysis Group at the Delft Institute for Applied Mathematics. In 2017, Mark Veraar was voted best teacher of the department of mathematics as well as of the entire EEMCS Faculty. He has been the runner-up three times.



Hard and soft techniques

Mathematicians like their models to be as close to reality as possible, making noise an important ingredient. Such models are often based on so-called *differential equations* (see frame). Not the easiest subject matter to wrap your head around, but as a former best teacher of the Faculty of EEMCS Veraar knows how to explain patiently and clearly. ‘A model and its solution are two entirely different things,’ he says. ‘It may be clear for an aircraft engineer that, when put in a wind

an indication of how efficiently you can make statements about a certain mathematical situation. It is often relatively easy to prove the existence of an optimum – a best possible outcome – using soft mathematics. Defining an efficient algorithm to find that optimal solution can, however, be very difficult. But sometimes it is much more difficult to use soft mathematics. In my research I try to stick with soft techniques, resorting to hard techniques only when absolutely necessary.’ At such a moment, a computer may be required

to perform a numerical approximation or a technical estimation.

A smooth solution

Thanks to the ever-increasing computational power, the mathematical models used by engineers and other scientists have become more and more realistic. ‘It may, however, still be the case that


Hard and soft mathematics

A soft mathematical approach is much more efficient for adding the numbers 1 to 100.

**Hard:** 1 + 2 + 3 + ... + 98 + 99 + 100 = 5050  
**Soft:** 1 + 2 + ... + 100 =  $\frac{1 + 2 + 3 + \dots + 50}{100 + 99 + 98 + \dots + 51} = 50 \times 101 = 5050$

tunnel, his experimental wing design results in a certain airflow. Using a mathematical model, however, it isn’t always possible to make any statements about that airflow.’ Veraar is interested in a fundamental understanding of the underlying equations, including noise. In his research, he applies both soft and hard mathematics (see frame). ‘Soft and hard has no relationship to the mathematics being simple or difficult,’ he explains. ‘Rather, it is

a prohibitively large number of calculations is required to approximate the solution of a model,’ Veraar says. ‘The better understanding we have of the properties of a solution, the more efficient an algorithm we can implement, thereby staying within computational limits.’ To attain such understanding, Veraar enthusiastically explores the depths of mathematics. One of the aspects he delves into is the so-called



‘It may be clear for an aircraft engineer that, when put in a wind tunnel, his experimental wing design results in a certain airflow. Using a mathematical model, however, it isn’t always possible to make any statements about that airflow.’



Partial Differential Equations (PDE)

A relatively simple example of a PDE is the heat equation. It describes the temperature distribution in a one-dimensional rod and how it changes over time.

$$\frac{\partial u(x,t)}{\partial t} = k \cdot \frac{\partial^2 u(x,t)}{\partial x^2}$$


The function  $u(x,t)$  describes the temperature for each coordinate  $x$  at time  $t$ . On the left of the equal sign you find the time derivative of  $u(x,t)$  – how quickly does the temperature at point  $x$  vary, in degrees per second. On the right side you find the second derivative to the location – if I look in the neighbourhood of  $x$ , do I observe an increase or a decrease in the speed with which the temperature changes. The heat equation is an exact description of the solution. It does, however, not give you the function  $u(x,t)$ , and thereby the solution.

*smoothness* of a solution. In mathematical terminology this is called ‘regularity’ and it is an indication for how often one can determine the derivative of a function (a measure for how quickly a function value varies, see frame), and whether or not these derivatives are ‘well-behaved’. ‘It is especially this smoothness of a solution that allows the implementation of efficient numerical approximations,’ Veraar explains.

Deep connections

His current research interests are at the cutting edge of three mathematical research areas (functional analysis, harmonic analysis and probability theory) and have been quite a treasure trove. ‘At a very fundamental level I have found unexpected connections between these research areas,’ he says. For a mathematician these are satisfying eureka-

moments while for an aircraft engineer it may mean that, next time, he may not have to build a scale model and put it in a wind tunnel. Nevertheless, Veraar is already considering a further expansion of his knowledge. His one-and-a-half-year-old daughter adds some very positive noise to his plans, but he may use the first semester of the upcoming academic year for an international research sabbatical. ‘I don’t yet know at what university I wish to broaden my perspective. There has to be sufficient overlap between their ongoing research and mine, so that we can understand each other. On the other hand, their interests should also be sufficiently dissimilar to allow me to write an interesting follow-up project proposal to my currently ongoing VIDI-grant.’ Wherever he may end up, noise will still be on his mind.



‘Noise can have a decisive influence on the predictive power of models’





# Using mathematics to explain why the River Ems is silting up

16

..... Text: Robert Visscher | Photography: Mark Prins and Herman Verheij

Why are the River Ems and the Dollard inlet so turbid that they barely harbour any life? Mathematician Henk Schuttelaars is working on new models to improve our understanding of this puzzling river and find solutions to the problems involved.

The River Ems and the Dollard inlet are located in the border area of northern Germany and the province of Groningen in the Netherlands. The water there has almost completely clouded up to the extent that it resembles liquid mud. Why is that?

It's caused by silt, a muddy substance that consists of clay particles and dead organisms. The amount of silt has been increasing over the past decades and parts of the river are basically now just liquid mud. This impedes the growth of algae, that are at the base of the food chain. Algae are the food of shrimps and shellfish, which in turn are eaten by birds. This has led to severe ecological degradation of the area and there is now almost no life left in the water.

*'I work on new mathematical models to better understand the enigmatic river and find solutions.'*

## But how did all that silt get there?

We're not exactly sure of the answer yet, which is why we are currently researching this. What we do know, is that much has changed over the past decades. For example, the Ems has been dredged to make it considerably deeper, for instance to cater for the increasingly larger cruise ships being built at the Papenburg shipyard in Germany. But land reclamation around the Dollard inlet has probably played a role too, because this has reduced the size of the estuary. We also want to know more about the origin of the silt and where it ends up. Before we can arrive at solutions to reduce the amount of silt and make the water clearer again, we first need to understand which processes affect this ecosystem. Otherwise there is no point changing anything; you might even end up making things worse.

## How are you conducting your research?

We're studying various factors that all influence the transporting of silt. One of these is the distribution of sediments throughout the water column, i.e. from the surface of the water to the bottom of the river. We also want to find out where the silt is entering the water and we are examining what happens under hyper-turbid conditions, which occur when a large amount of sediment collects in one place in the water. This influences the flow of the water and so changes the distribution of the silt. In short, we're studying a whole lot of things at once.





.....

In the past, among other things, water was not hyper turbid. If we know exactly what caused this, we can better predict the effect of future interventions.

**You are developing a model of the Dollard and the Ems to gain a better understanding of the processes at work. What kind of model is this exactly?**

Large numeric models are already being used by the Directorate-General for Public Works and Water Management and the Deltares applied research institute among others. However these models are not tailored to the situation in the Ems. As mathematicians, we are developing a new model that can be seen as complementary to the numeric models. One of the ways we are doing this is by looking into the past. The water was not hyper-turbid in 1965 and the area was ecologically healthy back then. We are curious to know what has happened since. If we know exactly what caused the turbidity, it will help us predict the effects of future interventions.

**How do you know what the situation was in 1965?**

It was anything but easy to find out. We didn't know if any data was available, so we just drove to the German city of Emden and tried our luck. We ended up in an archive in an attic with cupboards full of stacks of paper. It was a wintry day, there was no heating in the attic, and it was cold and draughty. We leafed through the papers, hoping to find something useful. And we did! We found tidal data from 1965; the numbers had been written by hand on a large sheet of paper. We also found soil maps with depth indications and samples of soil profiles from various years. We were very happy with this information, because it allowed us to reconstruct the situation in 1965. We discovered that in later years, and in particular from the 1980s onwards, the river became deeper and deeper and there was less underwater friction. The increase in silt made the river bottom muddier but also smoother. This resulted in less turbulence and other water movement, which was also reflected in the water levels. In 1965, the water level on the seaward side was higher than inland. It's now the other way round.

### About Henk

**Henk Schuttelaars** (1970) graduated Utrecht University in 1993 with a degree in physics. He obtained his doctorate at the same university in 1997 on the analysis of underwater soil patterns. He then participated in various projects at TU Delft and Utrecht University as a post-doc. Schuttelaars has been associate professor at TU Delft since 2006. He is currently building a new model to develop a better and more detailed picture of how silt behaves in the River Ems and the Dollard inlet. He hopes this will allow a better understanding of how silt has developed in the Ems estuary.



Photo: Herman Verheij





How did these discoveries improve your model?

We now have a better understanding of which processes played a role. Our model can describe what has changed between 1965, our baseline, and today. This will also tell us more about the role of turbulence and friction underwater during this period. When the river was made deeper, it also meant more sediment was introduced into the system. This decreased friction and turbulence which in turn led to more sedimentation. This was the cause of the extremely high sediment concentrations. There are other aspects, though, which we’re still unaware of, such as the influence of the sediment on the mixture in the water column, from the surface of the water to the bottom of the river, and how this changes under the influence of the tide.

In early September, a major survey was carried out involving no less than eight ships all across the Ems. How important was this?

One of the things the survey measured was the influence of the tide on the transport of silt. We can use this information to improve our model. It’s an interactive partnership; we can tell the survey team what information we need and then input the results into our model to improve it. That is why the September survey focussed on what happens to the silt in the river under the influence of the tide. Of course, the survey gathered a lot of other data too. We used special vessels to measure the water velocity, silt and salt concentration, and turbulence. Four ships were anchored in one place, while three others sailed back and forth across the flow of the river. One boat went from Papenburg in Germany and sailed downstream through the Ems estuary and into the Dollard inlet. This provided us with information about specific sites, but also about possible differences in the water over a long distance.

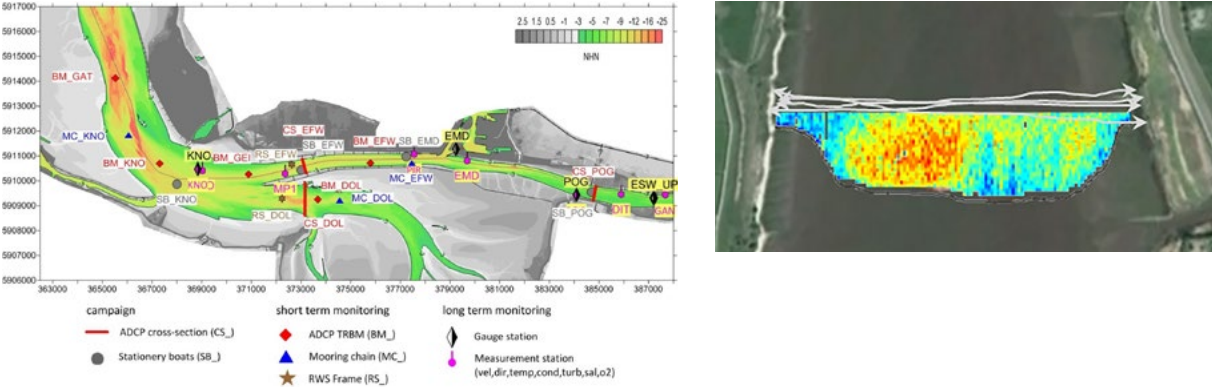
The survey involves a large number of parties mainly from the Netherlands and Germany. It is being coordinated by the Directorate-General for Public Works and Water Management and Royal Haskoning, but TU Delft, the universities of Oldenburg, Maine and Rostock (Warnemunde) and many local authorities are involved too. What is it like to work together with so many different parties?

We need all these partners, because otherwise we cannot produce a complete picture of the estuary. This is a huge project that encompasses a large region in two countries and involving an enormous number of measurements. That costs a lot of money, so you can only do it if you join forces. For example, we are using vessels provided by partners and highly specialised measurement equipment. We get together a few times a year to discuss progress and participate in workshops. I know almost everyone involved and we all work together well. We’re now waiting for the data from the most recent survey to be processed. We’ll import this data into our model too, and of course use it in publications. Similar areas such as the Scheldt estuary and some regions in Asia where heavy dredging takes place could also profit from our findings. It is in all our interests to cooperate, because it decreases costs and inconvenience, the river will be made safer and we’ll generate knowledge we can all use.


You were on board one of the research vessels during the survey. What was that like?

I sailed with one of the boats for three days. One of the measurements lasted 13 hours. We measured silt concentrations, among other things. I had to pull up a bottle filled with water once every twenty minutes. This gave me even more respect for the people who do this work on a daily basis. It is hard work on board. Ideally, I would have preferred to have collected data over a period of three months. It’s easy to think up a plan like that sitting at your desk just because you want the information, but it’s not really something you’d ask people to do, to work for so long on a project. I also enjoyed talking with the captain. While on board, I noticed that parts of the river flowed constantly while other parts did not. The captain explained that this is common and is caused, among other things, by small rivers branching off from the main flow. I was surprised to learn this. We spend all our time measuring and making calculations about the Ems, mainly searching for patterns, and so we view the river as a single system. By sailing with a survey vessel, I gained new insights and came to see the river in a new light.

'Henk is studying various factors that all influence the transporting of silt. One of these is the distribution of sediments throughout the water column; for example from the surface of the water to the bottom of the river. He also wants to find out where the silt is entering the water and we are examining what happens under hyper-turbid conditions, which occur when a large amount of sediment collects in one place in the water.'







# From micro to macro with duality

17

..... Text: Merel Engelsman | Photography: Mark Prins

Reality, how do we get there? We know how certain features of our macroscopic everyday world unfold and we capture these in mathematical formulas, but often these explanations do not have a basis in the microscopic world of colliding and randomly moving molecules. Frank Redig wants to understand the relation between microscopic and macroscopic reality, in particular for non-equilibrium systems. He was awarded an NWO TOP1-grant for his research into the conditions that must be met by a model of a non-equilibrium system, for it to be mathematically solvable using a powerful technique called *duality*.

*'Non-equilibrium means that the reversal of time results in a different movie being shown. For such systems, we try to mathematically explain how macroscopic laws emerge from the underlying microscopic world.'*

Mathematicians and physicists use models to describe reality and to try and understand it. Models are simplifications of microscopic reality, but sometimes they allow us to derive macroscopic laws of physics and to provide explanations for natural phenomena. For example, the freezing of water at a certain temperature or why certain materials can be turned into permanent magnets. These are, however, examples of systems in equilibrium. Models of non-equilibrium systems are often not mathematically solvable, and there is no general theory linking the micro-world to the macro-world, such as the Boltzmann-Gibbs distribution for systems in equilibrium.



## Behind the formula

'The physics of non-equilibrium phenomena is predominantly undiscovered territory,' says Frank Redig. 'We know what happens at equilibrium, and close to equilibrium, but not far away from it.' He explains that a metal rod, of which both ends have a consistent but different temperature, is a good example of a non-equilibrium system. 'We know the formula that describes heat conduction from the hot end of this rod to the cold end. This formula can be easily deduced from the law of the conservation of energy and the fact that heat flow is proportional to the temperature gradient. But this is a phenomenological description. What is lacking is a mathematical proof based on a model of the underlying microscopic world of individual vibrating atoms. I want to understand the emergence of the macro law from the microworld. How does the time evolution of the degrees of freedom of these individual atoms result in the heat flow? To develop such a proof, my colleagues and I apply statistical models that are simplifications of the microscopic quantum world.

## About Frank

**Frank Redig** completed his MSc and, in 1992, his PhD at the University of Antwerp. He subsequently spent six years as a monk and two years as a postdoc. From 2000 till 2005 he was assistant professor at the Eindhoven University of Technology, followed by four years of associate professorship at Leiden University and an appointment at the Radboud University Nijmegen. As of 2011 he is full professor in Applied Probability at the Department of Applied Mathematics of the Faculty of Electrical Engineering, Mathematics and Computer Science of the Delft University of Technology.

## Boltzmann-Gibbs distribution

The Boltzmann-Gibbs distribution is a probability distribution of the microstates of a system in thermal equilibrium. The higher the energy of a microstate, the lower the probability of this state to be occupied or present. The probability of a high-energy microstate increases with increasing overall temperature of the system. The Boltzmann-Gibbs distribution provides the exact compromise between maximum entropy (amount of chaos in a system) and minimal energy.

## Model

A model is a mathematical representation of a physical system. For example, particles jumping independently, and at random intervals, on a two-dimensional lattice. Or a chain of particles connected by springs to their nearest neighbours.





'To study these symmetries, the researchers use the so-called Lie-algebraic approach.'

### Many models, few solutions

'A prerequisite of the systems we investigate is that there is a conserved quantity. The total energy in the system, for example, or the number of particles,' says Redig. Even though the models are *caricatures*, and simplifications of the physical reality, most of them are not mathematically solvable. But if certain specific conditions are met, it may be possible to find an exact solution using duality. 'Most dualities so far have been found ad-hoc. Suddenly people found a much simpler system that could be connected to a certain non-equilibrium system. What we have developed is a constructive mathematical theory that can explain why certain models can be solved with duality and others cannot. It also tells us how to constructively find these systems and the corresponding duality functions.'

### Powerful symmetry

The mathematical formalism used by the researchers, is at a higher level of abstraction than duality itself, and depends on the presence of additional symmetries in the models. To study these symmetries, the researchers use the so-called *Lie-algebraic approach*. According to Redig 'this approach does not only indicate if duality is possible, it even automatically provides all duality functions. Depending on what you are interested in – the physical current, the number of particles, etc. – you only have to pick the appropriate duality function.' In a collaboration with Frank den Hollander from Leiden University and Cristian Giardinà from the University of Modena and Reggio Emilia (Italy), Redig received a TOP1-grant from NWO to further develop and apply this methodology.



.....  
In an infinitely large hourglass (sand timer), the same number of grains of sand will pass the narrow opening each second. If you reverse time, the grains will no longer fall downward but upward.

### Microscopic wealth

An important aspect of their research programme is the application of their theory to inhomogeneous systems, such as impurities in the metal rod mentioned earlier. 'Nearly everything in the world is inhomogeneous,' Redig explains. 'It only appears to be homogeneous when observed at a certain scale. Just like Redig's interests, the application of duality is not limited to fundamental physical processes. 'Think of the distribution of wealth in our society. The same mathematical models are applicable, because 'something' (money) is distributed while the total amount is conserved (if you look at the correct time scale). The fact that someone in the USA may handle money differently from someone in the Netherlands, can be modelled as a form of inhomogeneity. Through the ages, the distribution of wealth has satisfied the Pareto-distribution, 80% of wealth for the 20% most wealthy people. I want to know if a microscopic theoretical explanation exists, just as it does for the distribution of particle speeds in idealized gasses.

'We hope to discover equations describing a macroscopic system of neurons.'

#### Duality

Under certain conditions, a model that is not mathematically solvable may be 'linked' to a simpler model that typically has no physical meaning but that is solvable. The link between these two models is the so-called *duality function*.



#### Non-equilibrium

Non-equilibrium means that the movie of a physical system looks different if you reverse time. A flow is present. In an infinitely large hourglass (sand timer), the same number of grains of sand will pass the narrow opening each second. If you reverse time, the grains will no longer fall downward but upward.

### Dreamlike improvisation

Duality can also be applied to biological systems – the conserved quantity being the population size. Think of how population characteristics develop over time through evolutionary forces such as mutations, selection and migration. Though speculative, even human consciousness may be a consequence of a micro-to-macro transition, related to the physics phenomenon of a phase transition. 'A single water molecule does not freeze, but a large number of such molecules can undergo a phase change from liquid (water) to solid (ice). This phase transition can be mathematically explained by applying the thermodynamic limit – an unlimited number of particles in an unlimited volume. The same may be true for various phenomena in the brain. We hope to discover equations describing a macroscopic system of neurons.' This would mean that Redig, who thinks of both his mathematical work and his hobby of playing the piano as *dreamlike improvisation*, would possibly reduce his own creative processes to an equation.





## Affordable MRI

..... Text: Annelies de Bakker | Photography: Mark Prins and Io Cooman

Modern MRI scanners for disease detection are commonly used in the West but in developing countries hospitals are simply unable to afford them. Mathematician Martin van Gijzen is part of a team helping to develop a simple MRI scanner for a children's hospital in Uganda. 'Suddenly I am learning all about hydrocephalus.'

# 18






A simple MRI scanner can play a vital role in the diagnosis of hydrocephalus, a condition that affects over a hundred thousand children in sub-Saharan Africa. An early diagnosis can prevent severe brain damage. 'Once the outward signs of the disease manifest themselves it is usually too late,' Van Gijzen says.

Van Gijzen's field is numerical linear algebra, or the algorithms behind technical applications and simulations. He deals with numbers and abstractions and the results of his work are often not immediately obvious to the outside world. 'Writing a scientific article or developing an algorithm clearly matters to science but this project also has a much more direct social objective. It helps find a solution to a major problem in the developing world, one that involves children as well.'

### Weak magnet, complex algorithm

Van Gijzen goes on to explain his role: 'A normal MRI scanner has a superconducting magnet and liquid helium for cooling. It produces a complete, high-resolution picture. But scanners like this can cost up to three million euros and it takes highly specialised knowledge to operate them. Our aim is to develop a simple scanner which will cost no more than fifty thousand euros.' The reason it can be made so much more cheaply, Van Gijzen says, is down to the use of a weak magnet. 'This means the mathematical side of things is much more complicated and that is what we, that is what my colleague Rob Remis and PhD students Merel de Leeuw den Bouter and Xiujie Shan are working on.'



The joint venture of Rob Remis (Circuits & Systems), Martin van Gijzen (Numerical Analysis) and Andrew Webb (LUMC) is a great example of a fruitful collaboration. This video describes the partnership between Rob, Martin and Andrew.

*'This scanner is specifically used to detect hydrocephalus but the design, the magnet and the algorithms can also be used to diagnose other diseases.'*



Our aim is to develop a simple scanner which will cost no more than fifty thousand euros.

Using a weaker magnet makes the results of the scan less accurate. That means it is very important to gather as much information as possible via other methods,' Van Gijzen says. 'For example, better images can be made if we also use features indicative of hydrocephalus. By using all the data and combining images a simple scanner will, in the end, give us a fairly accurate reconstruction.'

### Multidisciplinary and intercontinental

Apart from the mathematical aspect the project is also an educational one for Van Gijzen. 'I come into contact with fields I am not very familiar with, medicine, for instance. And I am learning a lot about hydrocephalus, which is very interesting. At Delft I am also working with the DEMO engineers who are responsible for the hardware. So you see I'm not just staying in my own mathematical bubble.'

The partnership with British MRI physicist Andrew Webb, American neurosurgeon Steven Schiff and Ugandan biomedical engineer Johnes Obungoloch is a truly international one, which, Van Gijzen says, gives it something even more special. 'Andrew Webb works at LUMC and is a world specialist on MRI. Steven Schiff is head of the Neuroengineering Department of Pennsylvania State University and manages a large research programme on hydrocephalus in Uganda. Johnes Obungoloch is vice-dean of the new biomedical engineering department at Mbarara University in Uganda. A very diverse and inspiring group of people to work with!'



*'By using all the data and combining images a simple scanner will, in the end, give us a fairly accurate reconstruction.'*

### The team

From left to right: Merel de Leeuw den Bouter, Martin van Gijzen, Xiujie Shan, Rob Remis.



'At the start of the project, we were developing our first MRI scanner together with LUMC. That scanner never yielded images of sufficient quality. But we learned a lot from it. Based on the lessons learned, LUMC designed a new prototype that does give images. Together, we are improving their quality, both by image processing techniques and hardware improvements. We expect that this prototype will be ready for replication at MUST, our partner university in Uganda, in 2020.

### Contribution to the medical science in developing countries

Now the prototype is successful there will be a number of other challenges that will have to be overcome. The scanner needs to be able to withstand temperature changes and must be easily transportable to remote areas. The design also needs to take into account that babies tend to move their heads a lot. 'And of course it's important that doctors in Uganda can operate the scanner and that there are specialists on the ground to maintain it.'

Although the first phase of the project has taken place mainly in the Netherlands and the US, Van Gijzen says its ultimate goal is to broaden the base for medical science development in developing countries. 'The techniques and software will be publicly accessible. This scanner is specifically used to detect hydrocephalus but the design, the magnet and the algorithms can also be used to diagnose other diseases. In time the scanner will be manufactured locally. Most diseases are not that difficult to detect but you do need the right equipment.'

.....  
This scanner is specifically used to detect hydrocephalus but the design, the magnet and the algorithms can also be used to diagnose other diseases






## Organizing high-dimensional data with triangles and fields


..... Text: Merel Engelsman | Photography: Thierry Schut

Fast rendering of triangles appears to have little in common with neural networks. They are, however, linked by a clever insight by Nicola Pezzotti. Exploring the similarities in high-dimensional data through visualization no longer takes tens of minutes on a fast computer. The required computations can now be handled in real-time by a web browser. His breakthrough research is ideally suited for the validation of neural networks. But biomedical researchers at the LUMC are eager users as well.



'The mathematics of algorithms and scalability problems. That's what makes my heart beat faster,' says Nicola Pezzotti, former PhD-student in the Computer Graphics and Visualization group. He thinks of himself as an archetypical computer scientist, but the first ten minutes of our interview he highlights a number of practical applications of his work, with great enthusiasm. In a collaboration with the LUMC he discovered new cells of the human immune system. They also identified new differentiation stages of stem cells. 'Our data visualization indicated their existence. Knowing what to look for, the LUMC researchers were able to detect these cells.' The most important application of Nicola's research, however, is the validation of neural networks. Some background knowledge is required to understand his work on the t-SNE algorithm *t-distributed Stochastic Neighbor Embedding*.

*'The mathematics of algorithms and scalability problems. That's what makes my heart beat faster.'*





## Many dimensions

Nicola's work isn't so much focused on big data, as it is on high-dimensional data. 'This means that each data point has very many characteristics,' says Nicola. 'In the LUMC they can determine the presence of tens of different proteins on the surface of individual cells. That makes each cell a data point with tens of characteristics.' Visualizing these data would require tens of dimensions, something humans can't handle. Other examples of high-dimensional data are pictures of animals or hand-written numerals. When digitized, a computer can process and organize these. But with each pixel a variable (color, intensity, etc.), each data point has up to thousands of characteristics.

## Warping space

Whether it be the recognition of pictures of animals or the discovery of genes that may play a role in developing breast cancer, deep neural networks are the power tool to find similarities in such high-dimensional data. 'After a neural network has been setup for a particular application, it has to be trained,' Nicola explains. 'This requires a large set of training data. Pictures of animals, for example. The network will look for similarities that we know or suspect to exist.' The goal could be for the network to differentiate between a parrot, a pigeon and a giraffe. You tell the network what each training picture represents, and it will look for the characteristics that, apparently, fit this description. 'Colorful' can be one of the descriptive characteristics for 'parrot'. 'You can think of a neural network as warping high-dimensional space,' says Nicola. 'Pictures of the same animal will then form high-dimensional clusters, even if there is quite some distance between these data points in the unwarped space. When the network has been trained, it will be able to instantaneously label any new picture based on its location in high-dimensional space.'

*'A visual representation is an effective method to judge the quality of a neural network.'*

## About Nicola

**Nicola Pezzotti** is a researcher focusing on the development of algorithmic solutions. 'My algorithms and tools are used in several research groups and led to discoveries presented in important venues such as *Nature Communications*, *Nature Immunology* and the *Journal of Experimental Medicine*. I do not focus only on fundamental research, but also on the development of large, scalable and stable software and its path to useful products.' After his masters, he worked for more than three years in an Italian start-up developing 3D scanners. From September 2014 to October 2018, Nicola did his PhD at TU Delft. At the moment of publication of this volume, Nicola is currently working at Philips. 'As a research scientist in Artificial Intelligence, I lead the development of machine learning solutions for image formation in MRI. Among the other duties, I set up two research teams in collaborations with academic institutions.'

## An elephant is not a pigeon

It is, however, a bit more complicated. 'The output of a neural network is also high-dimensional,' Nicola explains. 'It is a series of numbers, a so-called *feature vector*, with each number a probability.' The first number may be the probability of the new image being a parrot, the second number the probability of it being a pigeon, and so on. All these probabilities add up to unity, to one hundred percent. 'If the network hasn't been properly designed or trained, it can't distinguish between the





*'It is essential that the low-dimensional representation preserves the characteristics of the high-dimensional data, such as their clustering.'*

various animals,' Nicola explains. The network may then label all parrots and pigeons, or all winged creatures, as one and the same animal. And if the pigeon is the only gray animal in the training set, it may label an elephant and rhinoceros as such. Developers also want to compare two or more neural networks. Which one is best-suited for the task? 'A visual representation is an effective method to judge the quality of a neural network, or to determine whether or not more training is needed,' says Nicola. Humans can only deal with up to three spatial dimensions. Therefore, plotting the (warped) high-dimensional space 'within' the network is not an option. But certain algorithms can provide insight into the similarities that exist in a high-dimensional data set, such as the output of a neural network. Nicola's research focuses on the t-SNE algorithm, which is especially well-suited for this task.

### Back to two dimensions

In 2008, Laurens van der Maaten - a former associate professor of TU Delft and now researcher at Facebook AI - published his t-SNE algorithm as an improvement of the existing SNE algorithm. The goal of both algorithms is to visualize the similarities in a high-dimensional data set by means of a representation in a reduced number of dimensions. 'This can be a three-dimensional representation, but in my work, I focus on the representation in a two-dimensional plane (2D),' says Nicola. 'It is essential that the low-dimensional representation preserves the characteristics of the high-dimensional data, such as their clustering.' Nicola explains that the t-SNE algorithm is a two-step process. Imagine a dataset of a few hundred thousand feature vectors, the result of feeding pictures of animals to a trained neural network. First, the algorithm determines the (high-dimensional) distances between all feature vectors. Data points that are alike, with little distance between them, will form a cluster in high-dimensional space. Second, the algorithm uses a process of optimization to search for the 2D-representation of all data points that best matches these high-dimensional distances. Data points that are alike should be clustering in 2D as well. Likewise, there should be large 2D distances between data points that are not alike in high-dimensional space.

### Neural networks versus t-SNE

Once training has been completed, a neural network can instantaneously classify new data. It can, for example, label a single subsequent picture to be 'a parrot'. The t-SNE algorithm visualizes the similarities within a high-dimensional dataset. Its output is a graphical representation of clusters of data points, without any indication whether a cluster represents, for example, 'birds' or pictures of 'a parrot'.

.....  
If the network hasn't been properly designed or trained, it can't distinguish between the various animals.



### Attractive and repulsive forces

The algorithm translates the distances between data points into attractive and repulsive forces. The more similar two data points are, the larger the attractive force. 'All data points are at least somewhat similar,' says Nicola. 'It is therefore important to have a repulsive force as well, or all points may end up forming a single blob.' The original SNE algorithm was slow because it calculated both forces between all data points. 'In this case, the number of calculations of distances and forces scales quadratically with the number of data points, as does the computing time,' Nicola explains. 'And you need many data points to get an intuition for the quality of the neural network.' It is this quadratic explosion in the number of calculations that makes the t-SNE algorithm a scalability problem, one to which Nicola is particularly drawn. Over the years, the original SNE algorithm has been improved a number of times. The t-SNE implementation by Laurens van der Maaten limited the calculation of the attractive forces to those points already in close proximity. 'The attractive forces are very small for data points that are unlike, that are far-removed. They can therefore be ignored,' Nicola explains. It was Nicola himself who recognized that the distances between the high-dimensional data points could be calculated with some inaccuracy. 'This dramatically reduced the computation time without sacrificing the end result.'

### Einstein's gaze

Nicola's most recent and breakthrough insight, however, relates to the computing of the repulsive forces. This was the only remaining time-consuming component of the algorithm. He perfected and implemented his insight during an internship at Google's Zurich office. His peers in visual analytics are more than enthusiastic. One fan twitters about 'new research from a clerk in our Bern office,' thereby referring to Einstein, who developed his theory of special relativity during his time at the Swiss Patent Office in Bern. A comparison with Einstein is typically flawed, but there certainly are similarities between Nicola's work and that of Einstein. 'Point to point interaction is the standard way to see the t-SNE optimization problem,' Nicola explains. 'You typically do not want to change the dominant paradigm. I decided to replace the point to point calculations for the repulsive force with the idea of a 'field'.'

In Nicola's approach, each data point in the 2D-representation projects a repulsive force onto a square piece of 'space,' thereby creating a small field. Other data points that are in close enough proximity will experience a repulsive force because of this field, and vice versa. Instead of calculating all these forces between data points one by one, Nicola sums up all small fields into a global field in the two-dimensional space.



At some locations these small fields may add up to a very strong global field, at other locations the global field may be much weaker or even be negligible. Subsequently, Nicola determines the repulsive force to each individual data point. This force depends only on the location of a data point within this global field. The field and the resultant forces are recomputed during each step in the optimization process. In its final implementation, the algorithm creates three global fields; for the horizontal repulsive forces, for the vertical repulsive forces, and a normalization field that ensures a smooth optimization.

‘It may seem that my algorithm performs unnecessary computations as it calculates fields at locations where there are no data points. But the fields have a very low resolution and, more importantly, the approach is ideally suited for implementation on a graphical processor.’ The square fields surrounding


each data point (see figure) consist of two textured triangles. Graphical processors have been optimized to process such triangles at a very high speed. As a result, Nicola’s implementation of the t-SNE algorithm is about a hundred times faster than any previous implementation. A researcher can now observe the optimization process in real-time, even within a web browser, rather than having to wait tens of minutes to see the end result. ‘My implementation can handle much larger datasets, and the real-time visualization provides additional opportunities. A researcher can interrupt the optimization process at any moment to decide on additional training of the neural network, or to zoom in on a specific cluster for additional detail.’

**Publicly available**

Nicola’s algorithm has been incorporated into *Tensorflow.js*, Google’s open source software platform for – among other things – designing, training and

analyzing neural networks. ‘It is important to me for software to be open source and readily usable by anyone. My software should not require a large computer infrastructure that only universities and companies can afford and maintain. I wanted my algorithm to be fast enough to be run on a single desktop computer.’ Nicola also published his research in scientific journals, but even that doesn’t satisfy him. For someone who is not an expert, analyzing high-dimensional data still is a very difficult task. ‘I am lucky to be part of a research group that consists of many excellent and highly motivated computer scientists, each with their own specific interests.’ His colleague Thomas Höllt is building a software platform, named Cytosplore, that bridges the gap between the algorithm and the biomedical researchers at the LUMC. ‘It is much more than a graphical user interface on top of some data processing routines. Thomas listens to their exact needs and creates a tool that allows them

to independently enter and analyze their research data.’ The applications of Nicola’s work are, however, not limited to neural networks and the biomedical field. A master’s student used his algorithm to analyze baseball player statistics. Nicola himself applied it to analyze social media filter bubbles. All that is needed now, is an algorithm that clarifies my electricity bill.




**Exploding complexity:**

The original SNE algorithm computed forces between all data points. Having three data points, there are three forces to compute; between the points A-B, A-C and B-C.  
Having four data points, there are six forces to compute (A-B, A-C, A-D; B-C, B-D; C-D).  
Having ten data points, there are 45 forces to compute.  
Having N data points, there are  $(N \cdot (N-1))/2 = (N^2 - N)/2$  forces to compute. The computational complexity scales with  $N^2$ .  
Having (only) 10,000 data points, there are nearly fifty million forces to compute!


For an acceptable calculation time, even for large data sets, computer scientists try to find solutions that scale with N or, at most, with  $N \cdot \log(N)$ .

Replacing forces with a field is very similar to what Einstein did. He replaced the gravitational force – between planets, stars and other objects – with curved spacetime. In his theory, the gravitational force experienced by any object is a result of the local curvature in spacetime.


Normalization Term



Horizontal Shift




Vertical Shift



.....  
A magnified representation of the three fields that each data point projects. The data point would be in the center of the three images. In the two images on the right, a negative value (blue color) represents a repulsive force directed towards the left or downward.





# So much more than 3D visualisation

..... Text: Merel Engelsman | Photography: Mark Prins


In recent months, EEMCS scientists Elmar Eisemann and Ruben Wiersma attracted a lot of attention with their three-dimensional augmented reality experience of a work by 17th century painter Pieter de Hooch. This is just the beginning of how mathematics and computer graphics will enhance our perception of art.

*‘There are similarities between the work of the Dutch Masters from the Golden Age and some of the research goals of our group.’*

‘I like to tell stories, using film or plain text,’ Ruben Wiersma says, ‘but the medium of computer graphics adds entirely new possibilities.’ The essence of a story is not so much in the beauty or realism of its setting, but in the message it conveys. It would therefore be too simple to characterise their 3D visualisation of the painting *Woman with a child in a pantry* as a gimmick, as a modern-day twist applied to a 17th century piece of art. ‘Pieter de Hooch paid a lot of attention to linear perspective to add depth to his paintings, just like his contemporary, Johannes Vermeer,’ Wiersma explains. ‘We specifically highlight this aspect of his work. We use computer visualisation to tell the story of De Hooch and his use of perspective.’

### The overlap between art and science


The idea for their visualisation is rooted in the organisation, by the Foundation Delft in-Ovatie, of an exhibition of the work by Pieter de Hooch, who resided in Delft from 1652 to 1661. Elmar Eisemann, professor in the Computer Graphics and Visualisation group at TU Delft, was one of the scientists approached by the foundation for a symposium to be held in parallel to the exhibition. He more or less succeeded in combining his love of mathematics with his childhood dream of becoming a Disney illustrator. ‘There are similarities between the work of the Dutch Masters from the Golden Age and some of the research goals of our group,’ Eisemann explains. ‘They were masters of light: the incidence of light, shadows, indirect lighting of surfaces, transparency, the accurate portrayal of complex materials such as clothing or metal objects. Within our research group, we aim to achieve a similar realism using computers, and to achieve this as efficiently as possible. Sometimes we have to develop completely new physics models.’





## How did the De Hooch visualisation come to be

First, each pixel was assigned a value for its depth, using an existing machine learning tool and the software for manual depth annotation as developed by the researchers themselves. They subsequently transformed this depth map into a three-dimensional model inspired by the perspective as used by De Hooch. The model had to be simplified as the augmented reality toolkit of the phone would otherwise not be able to provide a smooth experience. Using the phone's camera, this toolkit projects the three-dimensional model on top of (the image of) the painting, allowing the observer to move around the scene.



### Varying rules of perspective

De Hooch often painted the same scene twice. 'At first, our idea was to make a 3D reconstruction of a room using both paintings,' Eisemann says. 'That didn't work out well, as it turned out that De Hooch didn't always apply linear perspective very strictly. Sometimes he painted somewhat freely or at least applied different rules of perspective. We then decided to specifically focus on his use of perspective for adding depth. It fit well with our ongoing work in computer graphics in which we, for example, showed that you don't need to have maximum accuracy in depth for a scene to be perceived as being realistic. We had also already developed software to transform a two-dimensional image into an interactive three-dimensional scene, requiring only very little user input. That's how we came up with the idea of developing an augmented reality experience for the *De Hooch exhibition*.'

#### About Elmar

**Elmar Eisemann** is currently heading the TU Delft Computer Graphics and Visualization group, and he received the Dutch Prize for ICT Research 2019 for his research into the accurate, detailed depiction of visualisations using modern graphics hardware.

### Pinning down the perspective

The characteristic feature of linear perspective is that it uses one or more vanishing points located on a real or imaginary horizon (see figure). Lines that run parallel in the real world, such as the seams between tiles or the edges of a table, appear to recede to the



horizon towards the same vanishing point. 'To strictly apply this perspective, De Hooch marked such a vanishing point with a pin to which he attached a piece of string serving as a construction line,' Eisemann says. 'In some paintings you can uncover this tiny hole underneath the layers of paint using x-ray imaging techniques.' At first, the researchers tried to reconstruct these vanishing points by hand, just like many art historians still do. This turned out to be very time consuming as it required extreme precision. They therefore wrote an algorithm that automatically extracts lines and vanishing points from a painting. 'It turned out that, next to multiple vanishing points, De Hooch also used multiple horizons,' Eisemann explains. 'A room appears to be much larger when the central vanishing point of the floor is located below the 'real' horizon. It also increases the perceived distance

between the characters in the painting and you as an observer. We think De Hooch may have applied this false perspective intentionally.'

### A different perspective on art



The researchers also noticed that De Hooch often had his virtual horizons coincide with objects in the painting, such as the seat of a chair, a table top or a windowsill. The algorithm also exposed composition lines presumably used by the painter to arrange the scene and to direct the observer's gaze – a nobleman's pipe pointing at the eyes of the servant maid near the opposite edge of the painting, or a lower leg pointing at the pinnacle of a church tower. Quite often, De Hooch marked such composition lines with light and dark accents. 'All this information came to light using a very simple line detection algorithm, and it already tells us an intricate story of how the artist composed his paintings,' Eisemann explains. The augmented reality experience is only a simple application of the actual scientific research performed in the Computer Graphics and Visualization Group. Based on their reaction, the visitors of the exhibition and the symposium weren't any less impressed. 'Some visitors told us that they can no longer look at a painting



#### About Ruben

**Ruben Wiersma** recently started as a PhD candidate at the Computer Graphics and Visualization group. His interests include graphics, machine learning, and art. Ruben will start working on extending his MSc thesis research into geometric deep learning for publication and, in his PhD research, will work on computer graphics and machine learning applications for painting analysis under supervision of Prof. Dr. Elmar Eisemann and wProf. Dr. Joris Dik.





*'I find it interesting to apply mathematical precision in an area that, typically, is a bit less exact, a bit more elusive.'*

without wanting to find the imaginary horizons and vanishing points,' Eisemann says. 'And a photographer who didn't like the paintings by De Hooch, because he intuitively grasped the perspective to be off, now wondered how to use false perspective in his own work. Imagine that we employ the full power of artificial intelligence and computer visualisation to enrich the experience and understanding of art lovers.' That is exactly what Wiersma and Eisemann intend to do.

#### Combined with x-ray imaging techniques

Immediately after obtaining his cum laude master's degree, in October last year, Wiersma started with his doctoral research, again in the group of Eisemann. This research is partially financed by the Delft University Fund and


specifically focusses on the use of computer graphics for the analysis of paintings. 'I am a frequent visitor of museums and this research is a good fit to my general interest in art,' Wiersma says. 'We will use modern data analysis techniques, such as machine learning, to increase our understanding of paintings. Think of an even more profound analysis of the use of perspective, of the underlying layers of paint and of the colour palette used. The researchers teamed up with Professor Joris Dik, professor of Materials in Art and Archeology at the 3ME Faculty. Over the past decade he has used and improved x-ray techniques to non-destructively dissect paintings, layer by layer. Such an analysis results in massive data, which Wiersma intends to reduce to a more manageable

*'Imagine that we employ the full power of artificial intelligence and computer visualisation to enrich the experience and understanding of art lovers.'*

size using mathematical techniques. He will also apply advanced computer visualisation techniques to present his scientific findings to both experts and laypeople. 'The end goal of my doctoral research is the development of algorithms that can be used for any painting by any painter, and in any museum, such as the Mauritshuis and the Rijksmuseum.'

#### Art, not artificial

The question remains whether the research of Eisemann and Wiersma may somehow diminish the magic of art. 'Actually, we add magic,' Wiersma says. 'The presentation of art nowadays is very factual and scientific. They show you an x-ray image or a sample of the paint that was used. We can present this, and so much more, in an exciting manner, increasing the observer's connection to the piece of art.' Eisemann adds that this research is certainly not purely driven by the technical challenge. 'I find it interesting to apply mathematical precision in an area that, typically, is a bit less exact, a bit more elusive. Our De Hooch visualisation stimulated the observer to have a deeper experience of his use of perspective.' Both researchers stress the importance of having a continual discussion with the art community as to the use and added value of their research. Should you still have doubts about their use of computers to analyse paintings then, perhaps, Wiersma's personal opinion will reassure





The research into the use of perspective by Pieter de Hooch was a collaboration between Elmar Eisemann (professor), Ricardo Marroquim (assistant professor), Ruben Wiersma (doctoral student) and Yoann Coudert-Osmont (visiting student), all within the Computer Graphics and Visualisation group at TU Delft.

A three-dimensional visualisation of the painting *Woman with a child in a pantry* by Pieter de Hooch. In augmented reality, a computer adds objects and information to a live image of reality. In this case, the iPhone replaces the image of the painting by a three-dimensional representation. This representation updates in real-time whenever the phone changes its position or viewing angle relative to the image.

As an homage to the play with perspective of De Hooch, the application strikes a balance between a 3D reconstruction and the flat appearance of a painting.

you: 'I believe that art is communication between people. It is meaningful only because the person observing it has emotions, memories, experiences, that allow him or her to empathise. This is why art, by definition, is something uniquely human. A computer cannot diminish that.'





## How not to waste energy on 5G

21

..... Tekst: Merel Engelsman | Foto: Marcel Krijger

From an energy point of view, humanity's insatiable hunger for mobile data is unsustainable, unless you're willing to break a few long-held beliefs in the field of wireless communication. Professor Earl McCune was geared up to do just that.





**Many bits of energy make a big problem**

You may be jumping with joy when your mobile phone provider quadruples your data plan each year, for the next five years. A thousand times more data sounds great. 'But there is always a cost,' says Earl McCune, professor in the Electronic Circuits and Architectures group at TU Delft and specialized in sustainable wireless systems. 'All of the systems needed to provide these services, available on demand, draw power. The digital network core, the electronic systems processing and building the signals, the switching and data centres and, most of all, the transmitters and receivers.' Three percent of the global energy demand goes into wireless communication services. And with projections of the demand for mobile data to increase a hundred to ten thousand times over the next five to ten years, we are in trouble. 'The problem is that the communication industry in general has pretty much focused on trying to maximize the number of bits that can be transmitted per second,' McCune explains. 'Very little attention has been paid to the energy efficiency that goes along with that. I want to fix this.'

**About Earl**


On Wednesday 27 May 2020 our esteemed colleague **Earl McCune** has passed away. As of last year, Earl was (part-time) professor at the Electronic Circuits and Architectures research group at TU Delft and grew, in a short period of time, to become an incredibly popular expert in the field of sustainable wireless communication systems. In one of his last interviews (this article), Earl told us that he was willing to fight a few persistent beliefs in the field of wireless communication. Earl will not only be remembered for his boundless enthusiasm for the improvement of systems, and the inexhaustible amount of knowledge he brought with him; more than that, Earl was a mentor, an inspirer, an adventurer and a warm personality who put a smile on many people's faces with his dry sense of Californian humour.

**5G: a solution in the making**

Most of the hype in wireless communications focusses on 5G – the fifth-generation cellular network technology. 'The eventual switch to 5G comes with two major improvements,' McCune explains. 'One is the intended use of millimetre-wave directed beams for communication, saving lots of energy. The second improvement allows for a major increase in the speed of communications. This will enable, for example, internet of things controllable factories.' The official goal of the 5G communications standard is a ten thousand times increase in data traffic, at flat cost and flat energy use. 'What they didn't tell us, however, is how to exactly implement the standard and to achieve all this,' McCune continues. 'We know the math behind 5G, but it simply is not yet ready to be deployed on a massive scale. In order to reach the 5G energy efficiency goal, all links in the communication chain need to be overhauled.'

.....

**The official goal of the 5G communications standard is a ten thousand times increase in data traffic**







.....

The continuous effort to increase data rates has come with a continuous decrease in energy efficiency. Going from 2G to 4G, it has dropped from 60% to 20%

*'Presently rolled-out 5G networks are about as energy efficient as the old incandescent lightbulb.'*

### Power to the signal

The present paradigm in wireless communications, laid down in 1915, is to first build a very low-power signal. It consists of the addition of many sinewaves, simultaneously carrying a few bytes of information. Next, this jumbled signal is amplified to a communications useful power level of several watts. 'There is something called Ohm's law,' McCune explains. 'It says that if you have certain signal properties, then you will not be efficient at amplifying these, particularly when you require linearity in your circuitry.' Linearity means that the high-level output signal coming out of an amplifier is a faithful representation of the low-level input signal. To ensure linearity under all circumstances, the signal amplifiers will most of the time operate well below their maximum output level.

As a consequence, the continuous effort to increase data rates has come with a continuous decrease in energy efficiency. Going from 2G to 4G, the transmitter efficiency without any adjustment of its topology has dropped from 60% to 20% when going from 2G to 4G signals. A twenty percent efficiency means that for every watt of useful signal, four watts of energy will just end up as heat. 'For 5G, the efficiency will be only 10%, meaning that nine watts will be turned into heat,' says McCune. 'We're paying for the power, we're paying for the huge power supply, and we're paying to move that heat away into a giant block of aluminium serving as a heat sink.' For presently rolled-out 5G using millimetre-waves, which is still based on 4G like technologies, the efficiency is only a few percent. 'That's even worse than the old-fashioned incandescent light bulb,' McCune continues. 'The old tradition of using the linear amplifier, trying to find a version that is efficient, has run its course. But, in our digital age, physics gives us another way to make an accurate signal. If you have a sampling-based system, using square signals, Ohm's law then says you can be as efficient as you want it to be, bringing the efficiency of amplifiers up to 70%. That's a huge improvement, just by letting go of linearity and changing the circuitry. And we don't have to touch the signal standards themselves.'

*'In order to reach the 5G energy efficiency goal, all links in the communication chain need to be overhauled.'*

### Writing the textbook of next generation signal amplifiers

There is a catch, though. 'We need to build these square signals at communications-useful transmit power levels, not in the microwatt range,' McCune explains. 'There is no textbook explaining how to build a sampling system with high output power. We're going to be writing that.' Being both a part-time professor at TU Delft and part-time consultant and entrepreneur in California, McCune is involved in two separate approaches to build these untraditional amplifiers. In Silicon Valley, he pursues the idea of building one big switch, adjusting the power that is made available to it. Here at TU Delft, the focus is placed on a new solution consisting of hundreds of smaller transistors, in varying sizes. 'It's like playing a very big organ, adding the power signals up to make the desired output,' McCune explains. 'For the lower range of 5G frequencies, this has to be below 6 GHz. But most of the unused 5G bandwidth is the millimetre wave bands, having frequencies above 24 GHz,' McCune says. 'We don't yet have transistors with the capability that effectively can perform a switched operation at these operating frequencies, requiring transistor speeds in excess of 600GHz, so we need to improve them by another threefold or more.'

### No time to idle

It's not just the power supply for amplifying the wireless signals that has to be revisited. 'On average, when switched on, your desktop computer is not performing any useful tasks for 96% of the time,' says McCune. 'Likewise, the components making up the digital core of the communications network – the switching and data centres – are idle about 72% of the time. The logic components in the computer switch about a million times faster than the power supplies presently used. To have immediate performance when needed, we just leave the power on.' His proposed solution is quite like the modern cars that turn off and back on when waiting at a traffic light. It takes so little time that it doesn't interrupt your driving experience. 'We can already build the power supplies that respond in less than one microsecond,' McCune continues. 'The problem is that the computing chip community has interleaved their computer logic circuitry with some of the memory. Switch your computer off, and you'll erase a part of its memory too. We have to slightly rearchitect the chips, separating memory from logic. Then, we can reduce computing power consumption by a factor of four, saving millions of euros.'



## Line of sight communications

Probably the most hyped promise of 5G telecommunications is the use of directional beams for communication, transmitting power only where you need it. The present 4G network relies on antenna's that put out their radio signal across vast areas. As you move around campus, you'll always find the same signal everywhere. 'From a transmitter point of view, that is silly,' says McCune. 'This way, almost all radiated power is just going to warm up the grass. Using directed power, the antenna will create small beams, sending radio waves of the same strength only to where you are. When you move around, it will follow you like a theatre spotlight.' It works a bit like the satellite dishes used for satellite communications. But instead of parabolic antennas, it takes the coordination of hundreds of small antennas to create a beam covering only a very small segment of the sky. 'It is solved, but only in computer simulations,' says McCune. 'It will take some years for it to work in real-time in the real world, especially for 5G in the millimetre wave bands. The signal will be interrupted by bouncing off spinning windmills, cars driving by, the buildings people enter and exit. Even you yourself are a reflector, as millimetre-waves cannot penetrate the human skin.'

*'There is no textbook explaining how to build a sampling system at power. We're going to be writing that.'*

## The love for building things

One evening, when McCune was twelve, his now late father explained a few months' worth of network theory of electrical circuits to him. 'He was a good explainer, and it just made sense,' says McCune. He now has more than 45 years of experience in technology development in the field of radiofrequency and wireless design, most of which he spent in industry. 'I just love to build things,' he says, 'and I have seen a lot of what works, and what doesn't work.' Enjoying an early retirement in California, but still wanting to help solve society's problems, he accepted a professorship at TU Delft. 'When pressured, industry will change its technology only incrementally, just enough to keep making money. I thought we were missing something, there's so much room for big improvements.'

## The Fly's Eye project

'Now here's a challenge,' McCune says. 'Imagine the world cup soccer final with eighty thousand spectators who all, simultaneously, want to post their own high-resolution video of the winning goal. Can we do that?' This is what the Fly's Eye project is about, one of the ongoing TU Delft projects McCune is now involved in. 'It is very multi-disciplinary, with Nuria Llombart-Juan from the Terahertz Sensing group as the lead researcher,' he says. 'Instead of having spots moving around the stadium, the idea of the Fly's Eye is to 'just make a lot of spots,' covering the entire stadium. We simply hand the signal off to the next spot, should a user move around too much.' One of these passport-photo-sized spots is already being tested in the lab. It consists of a bunch of transmitters covered by a clear plastic lens that is able to focus radio waves in the millimetre-wave spectrum. The eventual structure will have 1500 of these spots, their beams partially



overlapping. At just over a meter in diameter, it will hang in the middle of the stadium. To handle the massive amount of wireless data, a bundle of optical fibres will connect the Fly's Eye to the internet, much like your optic nerves are the communication highway between your eyes and your brain. 'I just love the eye metaphor,' McCune says.

## Only a few minutes to dump satellite memory

McCune was also intrigued by the communications problem posed by cubesats. These tiny satellites, consisting of up to a few ten-centimetre cubes, are cheap to build and to bring into earth orbit. They are often equipped with one or more camera's for earth observation, collecting huge amounts of data. At their altitude of 500 kilometres they move so fast that direct line of sight communication with a ground station is severely limited. 'We have only a few minutes to transmit gigabytes of data, and there are some very limiting constraints,' McCune explains. 'There isn't much energy available in a ten-centimetre cube. The antenna can't be very big either, limiting us to high radio frequencies. Then, there's the oxygen in earth's atmosphere, which is very efficient at absorbing millimetre wave radiation.' His voice sounds more and more enthusiastic with every added challenge he mentions. 'It's a 5G type of problem,' he says. 'It involves energy, antennas, all kinds of things. To boost efficiency from two percent to fifty percent, we have to overhaul everything.'


.....  
Instead of having spots moving around the stadium, the idea of the Fly's Eye is to 'just make a lot of spots,' covering the entire stadium.



## Is sharing personal data a matter of all or nothing?

..... Text: Merijn van Nuland | Photography: Annelies van 't Hul

New digital technology is making it increasingly easy to share personal data. This brings all sorts of benefits, but it can also lead to loss of privacy and sense of responsibility. Birna van Riemsdijk, computer scientist at the Delft Data Science initiative at TU Delft and former lecturer of the Intimate Computing master course, is looking for the golden mean.



Technology is constantly on the move. Applications now available were unthinkable in the past, from sensors in watches that can count your steps to apps that help to improve your sleeping pattern. This technology is used not only to collect data for the user, but also to share this data with others.

### Keeping track

GPS trackers are tiny transmitters, often concealed in a watch which can, for example, be used to help parents track the location of their child. Very convenient, because it enables parents to keep an eye on what their child is up to. 'The problem is that with current technology you often have to share all the data at once. That might be done with the best intentions, but it also leads to all kinds of ethical concerns. Constant surveillance by a parent may affect the privacy of the child and lead to a diminished sense of responsibility. After all, it means there is always someone watching you who can intervene when necessary,' says Van Riemsdijk.

*'The problem is that with current technology you often have to share all the data at once. That might be done with the best intentions, but it also leads to all kinds of ethical concerns.'*



### Flexible technology

Van Riemsdijk is looking to find the golden mean with her research. She wants to take full advantage of the latest technology while doing justice to the privacy and responsibility of users. This will be done by making the technology's software more flexible. 'Indeed, how exactly data should be handled depends greatly on context,' says Van Riemsdijk. 'For example, you would probably like to keep a closer eye on a child with behavioural problems. The thing is that each parent has his or her own view on how to raise a child. There are wide differences among families, for instance, concerning the time of day children need to be at home. We are therefore aiming for technology that meets the norms and values of its users.'



*'Children considered the flexible version of the app more usable, and it gave them a greater sense of friendship and independence.'*

### About Birna

**Birna van Riemsdijk** is associate professor Intimate Computing at University of Twente. She is also guest researcher in the Interactive Intelligence section at TU Delft. Her research mission is to develop theory and software for creating intimate technologies that take into account our human vulnerability in supporting us in our daily lives: intimate computing is computing with vulnerability. For her research she was awarded a Vidi personal grant and the Dutch Prize for Research in ICT 2014. During this successful period she worked at the Faculty of EEMCS. 'At the Interactive Intelligence section at TU Delft the human-centered focus of my work really started to take shape. Not only by inspiring computational notions, or by studying how programmers use our created languages, but more and more through a focus on supportive technology: creating agents to support people in their daily lives, requiring collaboration between human and machine.'



### Integrating agreements

However, current digital technology is not yet capable of technically representing and using the various agreements made between users. This is going to change, if it is up to Van Riemsdijk. She is working on technology which will not only enable users in the near future to represent their mutual agreements on data sharing in the software, but will also enable the software to interpret these agreements. For example, the software could decide that agreements on safety issues take precedence over conflicting appointments of less importance.

*'We have to realise that technology is never value-free.'*

Developing such software is not straightforward. It means giving users much more freedom to define their data-sharing preferences. 'Of course, all those additional options improve the technology, but they also add to its complexity. It is therefore important to find out if users can still understand the technology.' All in all, the initial field study was quite encouraging!

Dr. Alex Kayal, one of Van Riemsdijk's graduates developed a new app for families for sharing location data as part of the COMMIT/ project, which was subsequently tested in an out-of-school day care centre. It turned out that the children considered the flexible version of the app more usable, and it gave them a greater sense of friendship and independence.

### Psychiatric disorders

Developing technology to share data in a more flexible way also has important applications in other areas, such as healthcare. Generating and sharing data could facilitate the work of healthcare providers, while patients can also use the technology to help manage a disease or disorder. The University of Twente and Scelta (GGNet), an expertise centre for treating personality disorders, are conducting research into the use of technology to treat people with a borderline personality disorder. There are indications that temper tantrums and self-injury or suicidal behaviour are preceded by physical signals such as an elevated heart rate. Once such signals can be identified, it may be possible to perform an intervention.

Together with Van Riemsdijk, research is being carried out into how this technology can be made more flexible with regard to data sharing between the therapist and the patient, based on the technology developed for GPS tracking. This too requires taking into account the privacy and responsibility of users.

### User freedom

Ultimately, Van Riemsdijk would like software and app users to have greater control of how they share their data, and with whom. 'We have to realise that technology is never value-free. Every new app or functionality reflects the designer's views on what it means to be human. It is important that we, the users, are able to determine the way in which we want to use an application. Sharing personal data should no longer be a matter of all or nothing. We must strive to find a golden mean. It is therefore important that we develop technology that supports this.'





# Solar cells available at bargain prices, but still the revolution stagnates

..... Text: Merijn van Nuland | Photography: Thierry Schut

In its Climate Agreement, the Dutch government aims to reduce CO<sub>2</sub> emissions by as much as 49% compared to 1990-levels by the year 2030. That's a laudable aim, says solar energy Professor Arno Smets, but achieving it will be a major challenge. 'It's difficult to estimate how much reduction in CO<sub>2</sub> emissions each measure will actually deliver.'

Arno Smets clearly sees the Climate Agreement as a positive step forward. He pays tribute to the fact that, with this agreement presented, the Dutch government is accepting its responsibility in the global challenge of the 21st century. He is also gratified to see the same government take the first far-reaching policy steps towards reducing the emission of greenhouse gases.

But there is still something troubling TU Delft's solar energy professor. Will it really be enough? 'I fear not,' says Smets. 'The energy transition is an extremely complicated political, administrative, economic, social and technological challenge for our society. Although some measures will definitely deliver better results than anticipated, achieving other key targets will not prove so easy or predictable in practice. That could quickly put the 49% reduction at risk.'

## Mission Innovation Champion

In the spring of 2019, Arno Smets has been named as one of the first Mission Innovation Champions. Mission Innovation Champions is a new initiative to honour pioneers in the field of clean energy. The programme is part of Mission Innovation, a coalition of countries that agreed to accelerate their efforts in the field of energy innovation at the Paris Climate Summit. Arno Smets was chosen because of the impact of his teaching and research on the energy transition. 'Climate change is the greatest threat to humanity. I believe we can develop sustainable energy systems that allow us to live in harmony with the planet. But we also need to change people's mind-sets and educate the world,' Smets says.

*'The energy transition is an extremely complex challenge for our society.'*



Unreliable energy sources

In his own research on solar energy, Smets has seen renewable energy production develop at an unimaginable pace. The production and upscaling of solar panels is happening much faster than everyone expects. Solar panels are being produced in Asia at bargain prices, making them available for less affluent households even in the Netherlands. And whereas the 2013 Energy Agreement did not feature even a single section on solar energy, in the 2019 Climate Agreement it has developed into one of the key pillars in the fight to tackle the climate issue.

‘Despite this, I can still see huge obstacles that we will need to surmount before solar energy can really become part of a sustainable energy supply,’ says Smets. ‘This is because wind and solar energy are still relatively unreliable sources of energy. A coal-fired power station is adjustable: you increase its output if you need more power. But with clean energy, you’re completely dependent on the wind and sun: when the wind drops and it’s cloudy, production is too low to meet the demand for energy.’


Tesla powered by grey energy

Another problem is the fact that electricity is difficult to store. Technically, it’s not yet really possible to store solar energy from the summer for the darker days of winter. This means we’re still largely dependent on grey energy for now, says Smets, which again has serious repercussions for other sustainability measures. ‘Take electric driving, for example. With government subsidies, you can certainly ensure that more Dutch people start driving electric cars in the short term, but what benefit will that have if all of the Teslas need to be charged using electricity from coal-fired power stations?’

Finally, but equally importantly, the integration of large-scale solar parks could be delayed. A solar park is a collection of a large number of solar panels that can quickly generate a lot of electricity locally. Smets: ‘In recent years, the government has been generous in offering subsidies for these parks, but the initiators struggle to find other investors. Grid operators can sometimes find it difficult to make the necessary expansions to the electricity network to keep up. What’s more, we’ll face a real shortage of electricians to install solar panels if this market continues to grow.’

About Arno

Arno Smets researches technologies to convert solar energy into usable energy, such as electrical or chemical energy. In 2009 he received a Vidi for his research into the use of thin silicon films to increase the efficiency of solar panels. In 2016, he was the first-ever winner of the edX Prize for Exceptional Contributions in Online Teaching and Learning. In 2015 he was appointed Antoni van Leeuwenhoek professor at TU Delft.



Arno Smets about the amount of energy Stephen Kiprotich, a Ugandan long-distance runner, produces during a marathon.

‘We are unaware of the complex world behind the wall socket.’







Take electric driving, for example. With government subsidies, you can certainly ensure that more Dutch people start driving electric cars in the short term, but what benefit will that have if all of the Teslas need to be charged using electricity from coal-fired power stations?

### Integrated approach

Smets' basic message is that solar energy is currently not yet a miracle cure. Existing scenarios may predict a key role for solar energy in the global energy transition, but all kinds of conditions will first need to be met for that to be achieved. According to Smets, the key is to pursue an integrated approach, because solar energy can only become a success with a good electricity grid, better opportunities for storage, sufficient investors and fitters, space for solar parks, social acceptance, etc. The success of any solar revolution in the Netherlands will depend on all of these factors. 'Things could progress quickly or slowly, which makes it difficult to make useful predictions. Currently, we just don't know how much CO<sub>2</sub> reduction every euro you invest will deliver.'

When he refers to an integrated approach, Smets means that all stakeholders take responsibility for achieving a smooth energy transition. An example: the last two years have seen a slight drop in worldwide investments in renewable energy. This lack of investment could prove an obstacle for the energy transition. 'I recently checked and my own pension fund, for example, invests only 1.3% of the budget in renewable energy. That's nowhere near enough to achieve all these projects.'

### MOOC's

Solar energy Professor **Arno Smets** has already reached hundreds of thousands of viewers with his MOOCs (Massive Open Online Courses). He has even inspired his viewers to install their own solar panels in places where there is poor electricity supply. Smets' online lectures are popular because of his lively style of presentation. But, of course, the topic also appeals to lots of viewers. 'When we started, no MOOC had been made about solar cells. We explain how solar power works and explore the fundamental principle of photovoltaic (PV) energy conversion. How do you convert solar energy into electricity? We demonstrate how the technology works and how you can build a modest system yourself. There's a lot of interest in sustainability and this appeals to lots of people.'

*'How is electricity produced, what is its impact on life on Earth and why isn't it done completely sustainably?'*

### Raising awareness

However uncertain the solar revolution may be, solar panels are good for one thing at least: raising awareness. 'All of us want to use electric devices and appliances, so we can use WhatsApp and watch Netflix. But we are unaware of the complex world behind the wall socket. How is electricity produced, what is its impact on life on Earth and why isn't it done completely sustainably? We don't ask ourselves these questions often enough. But if you have a few solar panels on the roof, this electricity production suddenly becomes much more tangible: you can actually reverse the dials in your meter cupboard!'

This helps consumers to understand the ethical issue that the energy transition really is. In that respect, the energy sector is quite similar to other sectors of industry, such as the clothing industry. We like to buy cheap T-shirts, but never wonder about the conditions in which they are manufactured. And we certainly don't want to see these poor conditions in our backyard, any more than we want to see wind turbines or large solar parks. But in a small and busy country like the Netherlands, there's no escaping the fact that the energy transition will make visible changes to our surroundings, feels Smets. If we want to make the transition to renewable energy, we need to be willing to make sacrifices.





# A post-coronavirus society

Text: Dave Boomkens | Photography: Marcel Krijger

The COVID-19 pandemic is not only putting society to the test, but also science. Scientists who have been able to spend their lives working quietly in the background suddenly find themselves the centre of attention. One such researcher is Piet Van Mieghem, a professor at TU Delft. He usually researches all kinds of complex networks that describe interactions. For example the internet or a more complex topic like the brain. But since March of this year, Van Mieghem has been using his knowledge as a weapon against COVID-19. At the request of the RIVM (National Institute for Public Health and the Environment), and with the help of three PhD students, he tries to predict the course of the virus and is working on a so-called exit strategy.

It was exactly ten years ago that Piet Van Mieghem began to investigate digital viruses that sneak into our networks via emails or system faults, with undesirable consequences. Now that we are experiencing the deadliest pandemic since the Spanish flu, Van Mieghem is shifting his focus from a digital to a biological virus. 'Because,' says Van Mieghem, 'the underlying mathematical models are basically the same.'

## NIPA

With this in mind, Van Mieghem and his colleagues are applying their special algorithm, NIPA\*, to the latest data supplied by the RIVM. They hope that NIPA will enable them to chart the progress of COVID-19. 'Epidemic models on any network rely on both virus specific parameters and the underlying contact graph, over which the virus propagates. If these parameters

are known, then the the spread of an epidemic can be computed. In particular, if the correct COVID-19 parameters and an accurate transmission map of the Netherlands were available, then we could determine the infection chance in any region.


You can calculate how many people are being infected per municipality or province. With NIPA we are doing exactly the same, but in reverse as we don't know the precise parameters associated with COVID-19. So instead, we are recording the number of infections per municipality per unit of time. The more frequently we sample the number of infected individuals, the more accurately we can determine the parameters for an epidemic model. Once we have recorded time series long enough, NIPA returns accurate enough parameters to predict the near future, i.e. a couple of days in advance.'

\*NIPA =  
Network-Inference-based Prediction  
Algorithm





*'Once we have recorded enough time series, we can start making simulations and models. So we are trying to look one, two, three or even four days ahead.'*



### Human brain

NIPA brings various academic disciplines together. One of these is network science, a relatively new discipline that studies the functions of a complex network and its associated underlying structure. We can use this discipline to study critical infrastructures, such as the internet and power grids, but also to study epidemics, or the human brain. Questions you may try to answer – when it comes to the brain – include: *what are the topological properties of such a brain-activity network? How does the structure of the brain change with age? Can you determine whether someone is intelligent (or has a high IQ) from their network?*

### Exit strategy


Back to COVID-19. Van Mieghem is also busy trying to propose and assess exit strategies. Together with a number of fellow professors, Van Mieghem is part of a team put together by RIVM to combat the coronavirus. For the time being, Piet also gets help from a PhD student who recently defended his doctorate, but is unable to return to China because of the coronavirus crisis. 'Since he has nowhere else to go, he is keen to do something useful and help us with this exit strategy,' explains Van Mieghem. 'For us this is actually a piece of cake, as we have carried out a good deal of research in recent years into the effects of earthquakes, floods and terrorist attacks on a network. This kind of thing puts a network under pressure, so it doesn't function perfectly everywhere. Questions we examined in recent years include *How quickly can you degrade a network?* and *What is the best strategy*

*for destroying a network?* And in this case we are also reversing things. The coronavirus crisis has degraded our society to a certain extent. What steps will we need to take together in order to return to some kind of normality soon? We call this *recoverability*. The Dutch government has imposed all kinds of measures concerning the coronavirus, but these cannot all be lifted at once as there are plenty of scenarios possible in which the virus could start to spread again.'

### Financial contribution

But that is still some way off. For the time being, Van Mieghem and his colleagues will be continuing to work overtime. In order to support Van Mieghem and all the other researchers at TU Delft who are working round the clock to combat COVID-19, we have set up the TU Delft COVID-19 Response Fund. The fund is being used to provide rapid financial support to researchers and students, ensuring no valuable time is wasted.





25

## Improving the effectiveness of speech recognition

..... Text: Robert Visscher | Photography: Mark Prins and Dave Boomkens

How can you ensure that a computer understands exactly what someone is saying even if they cannot speak the language perfectly? In her quest to make this possible, Odette Scharenborg is researching automatic speech recognition.

'Siri, find the quickest route home.' An instruction like that given in a quiet room is no problem at all for Apple's virtual assistant, used on mobile phones and tablets. But if you ask the same question in a busy restaurant or with a strong regional accent, Siri will find it slightly more difficult to understand what you said. This is not pure chance. Speech recognition software currently only works effectively if the person talking to it can be clearly understood, has reasonably standard pronunciation and there is little background noise. If not, we hit the limits of current technology, not only in the case of Siri, but also Google or Amazon, for example, who use similar software.

'Currently, speech recognition has little to offer to people with speech impediments, the elderly or someone who cannot speak clearly because of an illness. These are the people who could benefit most from it, for communicating with friends and family or controlling appliances and devices, for instance,' says Odette Scharenborg. As associate professor in TU Delft's Faculty of EEMCS, she is researching automatic speech recognition, exploring how you can teach a computer to recognise languages more effectively.

*'I actually knew nothing about speech, but at the very first lecture, I knew I was in the right place.'*





### About Odette

As associate professor in TU Delft's Faculty of EEMCS, **Odette Scharenborg** is researching automatic speech recognition, exploring how you can teach a computer to recognise languages more effectively.

### Psychology of language

The current technology works thanks to a huge body of data that is used to train speech recognition systems. The body of data is made up of two parts: speech and the textual description (verbatim) of what was said (the transcription). Recordings are made of the speech of people having conversations, reading texts aloud, speaking to a robot or computer or giving an interview or a lecture. All of these hours of audio material are written down and an indication is given of where there is background noise or where noises occur made by the speaker, such as laughing or licking their lips. But this system falls short. It is simply impossible (in terms of time, money and effort) to collect and transcribe enough data for all languages and all accents. Partly because of this, the focus has been on collecting data from the 'average person,' which means that the elderly and children are not properly represented. People with speech impediments can also find it difficult and tiring to speak into a machine over long periods. Equally, something said by someone with a foreign accent will often not be properly recognised because of a lack of training data.

There are also some unwritten languages that are still spoken widely. 'Take the Mbosi language in sub-Saharan Africa, for example, a language spoken by around 110,000 people. For the existing speech recognition systems, there is simply not enough training material available for that language. In addition, the main problem is that training a speech recognition system not only calls for the speech data, but also a textual transcription of what was said. Of course, you can record a lot, but quite aside from the fact that it takes a lot of time and costs money, transcribing what was said is not possible in an unwritten language,' says Scharenborg.

.....

In her current research, Scharenborg is combining insights from psycholinguistics with the technology of automatic speech recognition and deep learning.







On the ground floor of building 28 (Mathematics & Computer Science) there is a cabin that is about the size of a garden shed. In this cabin Odette Scharenborg's listening experiments take place.

Is it possible to develop a new form of speech recognition that does not use textual transcriptions, bringing an end to the exclusion of certain groups? This is what Scharenborg is currently researching. She is taking a fundamental look at automatic speech recognition. This involves such issues as how you recognise and categorise sounds. She is also exploring the psycholinguistic side, the psychology of language.

### New phonetic categories

In one of the methods Scharenborg is researching, she is looking at the sounds of languages and how they differ and what they have in common. Every language has its own set of sound classes, or phonetic categories, and associated pronunciation. Take the pronunciation of *bed*, *bet*, *bat* or *bad* in English. 'If a Dutch person pronounces these words, they will sound the same to many people. But that's not the case for an English person, for whom these are four different words with four different pronunciations,' says Scharenborg.

So even in languages like Dutch and English, which have a lot of overlap, there are different sounds (the vowel in *bed* and *bad*) and the same sound is sometimes pronounced differently (the *d* and *t* at the end of the word sound the same in Dutch but different in English). 'When developing a speech recognition system for an unwritten language, you can take advantage of the things that languages have in common. For example, the Dutch *e* sounds approximately like the English *e*, but the *a* in *bat* sounds something like a cross between the Dutch *a* and *e*. You can use that knowledge to train a model of the English *a* based on a model for the Dutch *e*,' explains Scharenborg.

In her current research, Scharenborg is combining insights from psycholinguistics with the technology of automatic speech recognition and deep learning.

Whenever we learn a new language, we create new phonetic categories. For example, when a Dutch person learns the English word *the*. This is different from the familiar Dutch *de* or *te* or the *a* in *bad* or *bat* mentioned earlier. The researchers aim to teach the computer this process of creating a new phonetic category. In the future, this should enable speech recognition to teach itself to learn languages in this way and recognise pronunciation, even if it deviates from the average.

### Safe door

The approach is based on experiments that Scharenborg conducted with human listeners. Test subjects listened to sounds or headphones. 'L and r are two separate phonetic categories, but you can also mix them artificially. The sound you then pronounce sounds just as much like an *r* as an *l*. We gave test subjects words that ended with the letter that sounds something like a cross between *r* and *l*. If you give them the word *peper/l*, they interpret the sound as a *r*, because it makes *peper*, which is actually a real Dutch word (pepper). After that, the test subjects start classing it as an *r*. The opposite happens if the same sound is given in the word *taker/l*, when people will interpret the same ambiguous sound as *l*, because *takel* is a Dutch word and *taker* isn't. You actually retune it in your head.'

'The current technology works thanks to a huge body of data that is used to train speech recognition systems.'

For these listening experiments, test subjects sit in a room that has been specially developed. On the ground floor of the Mathematics & Computer Science building, Scharenborg points to a cabin, roughly the size of a garden shed. She opens the heavy, wide door. It is so thick that it looks like a safe door. 'It's in here that we do the experiments. There's no distraction from other sounds. Can you feel the pressure on your ears? That's because the room is low-noise,' says Scharenborg.

In the cabin, there is a chair and a desk. On that, there are headphones and a computer. On the screen, test subjects can read words and they can hear sounds via the headphones. There is also a window. This ensures that the testers always have contact with a supervisor who can see if everything is working as it should.



The data that the scientists collect in the test is being used to make a new model that should ultimately result in better speech recognition. In it, the model itself judges whether or not a new phonetic category is needed and you no longer have to enter it. This is a major step forward for automatic speech recognition. The model was inspired by our brain and our deep neurological system, says Scharenborg. It is known as a deep neural network. 'We are the first to investigate self-learning speech recognition systems and we are currently working with a colleague in the United States to develop technology for this. I expect to see the first results within two years.'

#### Bridging the divide between humanities and sciences

Will this soon produce an effective system that can rival Siri? We cannot really expect that any time soon. The technology is still in its very early days. Scharenborg's research is primarily fundamental in nature. Despite that, it is laying the foundations for a new type of automatic speech recognition. A program that works perfectly is still a distant prospect.

In a previous NWO Vidi research project, Scharenborg also explored speech recognition. 'Everyone knows that it's harder to understand someone you are talking to if there's a lot of background noise. It becomes even more difficult if someone isn't speaking their native language, but another language because they want to order something to drink in a busy café on holiday. But why is it so much more difficult to understand each other when there's background noise? And why is this more difficult in a foreign language than in your native tongue? I was keen to find that out and ended up learning more about how humans recognise speech.'

In her current research, Scharenborg is combining insights from psycholinguistics with the technology of automatic speech recognition and deep learning. Normally, these are separate fields. Psycholinguistics is mainly the preserve of scientists and social scientists and focuses on language and speech and how our brains recognise words. Speech recognition and deep learning are about technology and how you can use it to make a self-learning system.

'In order to make further progress, we need to combine both fields. It seems rather strange to me that there are not more people who combine them. The aim of the people and the machines is the same: to recognise words,' says Scharenborg.

*'Everyone knows that it's harder to understand someone you are talking to if there's a lot of background noise.'*

The current technology works thanks to a huge body of data that is used to train speech recognition systems.



The academics working in the two different fields often don't even know each other. This is something that Scharenborg aims to change. For years, she has been organising the largest conference on the subject of speech recognition, bringing together both communities. 'Both groups still focus too much on their own field. For applied scientists working on deep learning, speech is a type of data that they use to develop technology. However, psycholinguistic researchers often have relatively little knowledge of computers. In my view, these fields complement and can learn from each other, in terms of knowledge about both speech and computer systems. I want them to engage in discussion.'

#### Knight Rider

Scharenborg has extensive experience of both automatic speech recognition and deep learning and therefore aims to bridge the divide. Before joining TU Delft, she was associate professor in the Linguistics department at Radboud University. 'My own field is speech recognition, but I also work with psycholinguistics. It makes logical sense for me to do my research at TU Delft now, because the technical expertise is here and research is conducted across disciplines. That's why I don't feel like an outsider, even though I'm virtually the only person in the world doing this research.'

Even at secondary school, Scharenborg was good at both humanities and science subjects. 'I was unsure about what to study, which is why I chose a programme focusing on language, speech and computer science at Radboud University. I actually knew nothing about speech, but at the very first lecture, I knew I was in the right place. Almost everyone talks to each other and a lot of communication is done through speech. You can look at it from a physics perspective by focusing on waves that propagate speech, but also explore how people recognise it. Being able to approach it from different sides really appealed to me.'

Without actually realising it, she was already working with speech. 'For example, I really loved the TV series Knight Rider. In it, David Hasselhoff played the hero Michael Knight, who could speak to his car KITT. This supercar not only understood what was said, but also talked back and was intelligent. Knight and KITT also used a wrist watch to communicate when they were apart. Even then, I found it really interesting and that fascination has never gone away – I'm still working with speech and technology.'





## Diagnosing tuberculosis with algorithms

Text: Koen Scheerders | Photography: Mark Prins

We've been able to decipher DNA for decades now. Until recently, however, the information behind this blueprint of life was inaccessible. Now mathematical algorithms unlock this information for us. And that has highly promising applications, such as diagnosing complex tuberculosis infections.

26



The abstract world of mathematics is far removed from a bed in a hospital in a developing country. Take, for instance, the Church of Scotland Hospital in Tugela Ferry, a small country town in South Africa. The hospital was unknown to the rest of the world until it experienced an outbreak of tuberculosis, a contagious bacterial disease, in 2006. This extensively drug-resistant tuberculosis (XDR-TB) proved very difficult to treat and spelt death for most of the patients in the hospital. In almost no time, more than 300 patients with an XDR-TB infection were reported throughout South Africa. For lack of proper treatment, the majority of these patients died.

*'The idea is **sequencing for anybody, anywhere**. Combined with our algorithms, this gives us a test that tells us it's time to see a doctor.'*

**Stubborn**

Now, more than ten years later, this extensively drug-resistant family of TB bacteria is still responsible for a quarter of all XDR-TB cases in the area round Tugela Ferry. The strain has become an example of the major problem that resistance to TB anti-

biotics presents in developing countries in Africa. Nearly everyone there carries this bacterial infection, which is easily spread from person to person through the air. If a patient goes to hospital for treatment with antibiotics, the doctor cannot see what form of TB the person has. Treatment is therefore completely effective in only 20 per cent of all cases: the person is often prescribed either too much, or not enough, antibiotics. This inadequate treatment can lead to resistance: TB strains that 'get used' to drugs and for which increasingly specific antibiotics are needed. This can lead to an infection that is increasingly difficult to control, such as with MDR-TB, which can easily spread within South Africa and beyond.

**Complex infections**

In Delft, Thomas Abeel is working on a solution to this global problem. Together with partners, such as hospitals in South Africa, he diagnoses cases including complex TB infections, which consist of multiple types of TB bacteria. But instead of a blood test or X-rays, he is using algorithms. 'Complex infections are therefore more difficult to treat, which increases the risk of contagion,' Abeel explains. 'One of our findings is that there are many more complex infections

.....


Thomas Abeel's algorithms reveal resistance patterns in TB strains and create prospects for new treatment options.

than we thought. Each strain of bacteria has its own resistance profile. And a combination of strains can lead to a resistance that we are no longer able to treat.' Such as the XDR-TB infection that startled South Africa. Since then, 10 to 30 per cent of TB infections are classified as complex.

**Resistance patterns**

Using patients' genetic data, the algorithms that Abeel has come up with can *pull apart* such a complex infection at genetic level and derive which kinds of TB the patient has in his body. Abeel: 'An algorithm is a kind of computer program that is based on a model. Like the principle that a sequence of DNA is a combination of DNA from multiple sources. We can turn that into computer code. That enables us to divide the sequence of tuberculosis DNA and find out what the original TB strains were.' And that is greatly needed, because the number of patients with complex TB infections continues to rise. Abeel: 'The biggest problem is that 99 per cent of the infections occur in developing countries. Those are the countries with the greatest resistance to antibiotics, but they are also the countries with limited funds for diagnosis and therapy.'

*'Each strain of bacteria has its own resistance profile. And a combination of strains can lead to a resistance that we are no longer able to treat'*







*'There's no reason to assume that we cannot apply our algorithms in other fields.'*

### Chemotherapy

Thomas Abeel's algorithms reveal resistance patterns in TB strains and create prospects for new treatment options. It takes more than just a couple of weeks to treat someone for TB; it is a long and arduous treatment. Abeel: 'People often take four kinds of antibiotics for two months, followed by two kinds for another four months. For an MDR-TB infection, the regime lasts two years and is often much more complex. The drugs have serious side effects, such as skin discolouration and disorientation. It's almost like chemotherapy.' Many patients often feel much better after a week and therefore stop that arduous treatment too early. And that breeds even more resistance. Abeel: 'If we know exactly which antibiotics still work for each strain, we can diagnose and treat complex TB patients more precisely.' That increases the chance of being cured immensely.

### About Thomas

**Thomas Abeel** studied Computer Science Engineering in Ghent. After graduation, he took a PhD in Plant Biotechnology, followed by a three-year internship at the Broad Institute in Boston, part of the Massachusetts Institute of Technology (MIT). It was there that he started his work on the genetic analysis of tuberculosis. Since 2014, he has been recently promoted to associate professor of Bioinformatics in the Intelligent Systems Department of the Faculty of Electrical Engineering, Mathematics and Computer Science at TU Delft. He still holds his guest appointment at the Broad Institute.


### Technology

Nevertheless, Abeel says we may still have to wait a long time before complex TB patients can be treated individually. One of the reasons lies in the technological challenges. For instance, how do you obtain genetic material in a developing country? Abeel: 'We have all the components, but it will still take several years of technological development before everything can really all be used together.' One of the most promising developments in this area is nanopore sequencing: unravelling patients' DNA with a small device that looks like a USB stick that fits in your laptop computer. The company that is developing this is also working on a DNA pen that extracts DNA directly from blood or saliva. The rear of the extraction pen emits a drop of prepared material that is fully suitable for nanopore sequencing. Abeel: 'The idea is 'sequencing for anybody, anywhere'. Combined with our algorithms, this technology gives us a test that tells us: yes, you have TB; time to see a doctor. Or that a doctor can use it to make a diagnosis and a suitable treatment plan.'

### Application

But it doesn't stop there, Abeel thinks. His algorithms can be used in an even broader context, such as by applying them to the food industry. 'Maybe we want to be able to check the chicken in the supermarket for resistant bacteria? Or food producers might want to check their yoghurt cultures or strains of yeast for contamination? DNA sequences form the basis, so there is no reason whatsoever to assume that these algorithms cannot be used elsewhere. That, too, is diagnostics.'





# 27

## Tackling tumours using mathematics

..... Text: Dave Boomkens | Photography: Mark Prins

Radiotherapy is one of the treatments for cancer, but designing a treatment plan is a challenging task. What is the best way to distribute the unavoidable high doses to healthy surrounding tissue, as to avoid complications? Mathematical models have an important role to play in ensuring that radiotherapy can be conducted with minimal impact on the patient's quality of life. Marleen Keijzer (lecturer at the Delft Institute of Applied Mathematics) and her former EEMCS student Sebastiaan Breedveld (assistant professor at Erasmus MC) discuss the role of mathematics in an ever-changing field: radiotherapy.

### Marleen Keijzer

#### Marleen, what first sparked your interest in mathematics within radiotherapy?

In the mid-1980s, I attended a conference in America on how laser light is used for medical applications. For example, in the case of atherosclerosis, how do you deal with a constricted blood vessel? And how could lasers be used to treat a particular type of cancer? At the time, this was all relatively uncharted research territory. Soon after, I moved from Delft to the United States in order to apply a Monte Carlo method – in other words mathematics – to simulate light scattering. I wrote articles and worked with various hospitals, but decided to return in 1989. Back in the Netherlands, I completed my PhD thesis on calculating light propagation in tissue. If you put a flashlight against your cheek, most of your face will turn red. That is scattered light. But how does the light actually get there? After my PhD, I worked on several more research projects, but soon realised that I much preferred teaching and supervising students with their graduation research. For the latter, the focus was actually always on the useful application of mathematics, for example within radiotherapy.

*'If you put a flashlight against your cheek, most of your face will turn red. That is scattered light. But how does the light actually get there?'*



Final-year students and radiotherapy. What challenges lay ahead of them?

In the field of radiotherapy, it was still possible to make major strides. Graduating students were given all the space needed to design and program their own optimisation methods. An optimised plan of this kind gives you the maximum opportunity to control the tumour and minimise the risk of damage to healthy tissue. So you search for the desired dose distribution and use optimisation to determine how the beams (incidence directions from which the patient is irradiated) should relate to each other. What happens if you administer one beam here? And what consequences will that have for the dose applied there? Not so long ago, this still had to be done manually: a radiotherapy technician set the beams, an advanced computer program calculated the dose distribution and the radiotherapy technician then adjusted the beams to achieve the optimum dose distribution. The problem with radiotherapy is that it doesn't just involve one single criterion, but numerous conflicting criteria. Optimising a radiotherapy plan could therefore easily take up a whole day.

Sebastiaan Breedveld was one of your graduating students. What exactly did he do?

I supervised Sebastiaan for a while, but he eventually did his PhD research at Erasmus MC – Daniel den Hoed. He was awarded his PhD cum laude on the subject of optimisation. Because radiotherapy technicians were having to devote a whole day to a radiotherapy plan, Sebastiaan developed a method in which optimisation was fully managed by a computer. An important component of this is the so-called *wish list*. For some organs, the radiation tolerances are well known. You're dealing with strict limits. It's important to take account of these limits, because if the radiation dose is too high, for example in the case of spinal treatment, it may leave the patient paralysed. But the reverse is also true: if you radiate a tumour using less than the minimum dose, it may start to grow back. As well as adhering to strict limits, you need to keep the dose administered to each organ as low as possible. The wish list allows you to do this in a prioritised way (one by one). In order to guarantee quality of life, one organ is sometimes more important than the other. This is something you will also take account of when developing a treatment plan.

About Marleen and Sebastiaan

**Marleen Keijzer** is a lecturer at the Delft Institute of Applied Mathematics, also known as DIAM, at TU Delft. As part of her job, Keijzer is responsible for the first-year course in Mathematical Modelling, and teaching and coaching students is what she enjoys most. In the last decades, Marleen has supervised numerous students during their graduation projects or PhD programmes. Most of her projects were based at the Daniel den Hoed Clinic – now Erasmus MC – and focussed on radiotherapy treatment planning. Together with her colleague Theresia van Essen, Marleen is currently setting up an online course In Hands-on Optimisation. Keijzer also recently became vice-chair of the TU Delft Works Council.

**Sebastiaan Breedveld** studied Applied Mathematics at TU Delft. Both his Master's and PhD research focused on the automation of treatment plans within radiotherapy. Breedveld is an assistant professor at Erasmus MC, where he is continuing his research aimed at improving radiotherapy treatment methods.







### So this radiation issue is really about optimisation...

Yes, that's actually what it comes down to. After all, numerous optimisation methods are being used in radiotherapy. Nowadays, the maths are done by computers. But it still takes time. The graduation students I am now supervising are working to make this process faster and more practical while also tailoring it more effectively to the individual patient. Ideally, you want to fine-tune the plan when the patient is actually lying on the treatment couch of the radiotherapy device. Real time. The human body is always moving, while a radiotherapy plan is made weeks before the start of the first treatment. If the bladder is slightly fuller, it may have pushed an organ or tumour into a slightly different position, for example. With that issue in mind, we had another doctoral candidate develop a statistical model of the movements of the prostate. He analysed numerous CT scans from different patients, identified the most important movements from these and summarised

it all in a model. This is a very useful application for mathematics, because as soon as you realise where the uncertainties are, you can start taking them into account. It is a very practical EEMCS issue.

### What challenges are there for the future?

Radiotherapy techniques could be customised more. In the 1960s, they would put a piece of lead on a patient, with an opening cut out for the radiation beam at the level of the tumour. This meant that the radiation beam was limited to the tumour alone. Today, everything is done using three-dimensional CT scans, fast computer calculations and highly targeted radiation beams. There has been significant progress, but treatment could still be tailored more effectively to the individual patient. Also, planning methods are still not anything but standard. It would be great if planning methods were generally accepted and made easily accessible, so their use would not be limited to the university hospitals of wealthy countries.

Pictured is a radiotherapy treatment unit. The extension at the top-left can rotate around the patient, and deliver beams of ionising radiation. By rotating to a different position, the patient is irradiated by another 'beam' direction. The number of beam directions, their directions, and the intensity distribution within each beam are degrees of freedom which allow the delivery of complex dose distributions inside the patient. Obtaining optimal settings is our challenge, and we use mathematical optimisation to obtain them.

## Sebastiaan Breedveld

### Sebastiaan, what first sparked your interest in mathematics within radiotherapy?

I was always eager to do something in the world of medicine. Eventually I opted for a study in Applied Mathematics at TU Delft. But I kept being attracted by the medical field. With that attraction in mind, I decided to do my graduation research on a mathematical issue with a medical application. Because Marleen lectures on the subject, I thought: 'I would like her to supervise my graduation research!' Via Marleen, I ended up at Erasmus MC – Daniel den Hoed in Rotterdam.

### How did you manage to combine mathematics and radiotherapy there?

Maybe I should tell you something about radiotherapy first. Radiotherapy is used in approximately half of all diagnosed cases of cancer. Radiotherapy is all about irradiating tumours – you're actually destroying cells. To do that effectively, you need a radiotherapy treatment plan. This plan actually describes the settings for the radiotherapy device that will result in the desired dose distribution. What makes it problematic is that you're dealing with ionising beams that go straight through a patient. Such a beam actually damages everything in its path. On the one hand, you of course want to apply enough radiation to be able to eliminate the tumour. But then again, you don't want to do this at the expense of the surrounding organs: your aim is to minimise the risk of complications as much as possible. Take a tumour in the head and neck area. If you expose the small salivary glands to a dose above its tolerance during irradiation, this may result in permanent damage. It may lead to a situation where the patient has to drink a little water every 30 minutes, even at night. What dose do you opt for? And how do you then distribute that dose? These are important questions that involve mathematical calculations.

*'In radiotherapy, together with each problem solved, a basis is created for new opportunities to further enhance treatment!'*

Radiation therapy is one of the most important treatments against cancer. One of the ways to do this is through external beam radiotherapy (EBRT). A Linac (a linear particle accelerator) is often used for this treatment. The position of the angles is one of the problems you have to deal with during the preparation of the treatment. In other words: from which directions the patient is going to be irradiated. As you can see in the video, there are different degrees of freedom that can be used and combined. The device in the video is used in the majority of the hospitals. For both traditional and more contemporary treatments, think of proton therapy.

SCAN!





How did you include these questions in your own research?

It's worth knowing that my research involved two stages. The first stage was to investigate how the available equipment could be used as effectively as possible. In the Netherlands, we've seen a rapid development of most equipment, while the range of possible applications has lagged behind. Smart use of mathematical models enabled us to catch up and put the available equipment to maximum use. The second step was to answer the question: how do you distribute the dose? In order to find out, I asked a number of colleagues: what is it you would like to see in a therapy plan? Based on that input, I compiled a wish list: a list that I could use to guide the optimisation. Running through this wish list, you can work towards a radiation dose that is highly precise based on sophisticated calculations. At that time, the process was relatively new.

*'For some organs, it's perfectly clear how high the radiation dose can still be. You're dealing with strict limits.'*

How did you apply the wish list in practice?

This wish list ultimately generated automated working methods. Now, there were two plans: one plan that had been developed clinically and one plan generated automatically. In order to find out which method worked best, both a clinical and an automatically generated plan were developed for 50 patients. By doing this, we hoped to achieve some kind of statistical relevance. Both variants were submitted to the treating physician who was then asked: which plan would you use to treat this patient? The doctor had not been told which variant had been calculated manually or automatically. After repeating this 33 times, we ended the experiment. We found that in 32 out of 33 cases, the preference was for – what turned out to be – the automatically generated plan. The difference in quality was so convincing that from an ethical perspective we felt compelled to make an effort to treat all patients within this group in the new way. In other words, the wish list worked.

.....  
What makes radiotherapy interesting is that new techniques are continually emerging. For example, proton therapy has been around for a decade, but at Holland PTC, the very first patient was only treated December 2018. In radiotherapy, together with each problem solved, a basis is created for new opportunities to further enhance treatment!

In the years ahead, what will be the greatest challenge in the field of radiotherapy?

The HollandPTC outpatient centre – based on the TU Delft campus – recently started treating patients using proton therapy. This form of radiotherapy for cancer is new to the Netherlands. Protons are small, positively charged atomic particles that enable high-precision irradiation. This means that less radiation is delivered to healthy tissues, thus reduces the risk of side-effects. TU Delft is working on the technical side of this study. Rotterdam and Leiden manage the medical side. But there's another question: how can you generate the radiotherapy treatment plan even faster and more effective? Imagine you need to make last-minute changes to a plan. How then do you ensure you can go through all those gigabytes of active data in just a few seconds?


*'The problem with radiotherapy is that it doesn't just involve one single criterion, but numerous conflicting criteria.'*

Would you ever change the world of radiotherapy for another one?

When I started here, I thought: when I've done my bit, I'll slip away and move into hydraulic engineering. I've now passed that point. What makes radiotherapy interesting is that new techniques are continually emerging. For example, proton therapy has been around for a decade, but at Holland PTC, the very first patient was only treated December 2018. In radiotherapy, together with each problem solved, a basis is created for new opportunities to further enhance treatment!







# Ultrasound from a plaster

..... Text: Koen Scheerders | Photo: Mark Prins

Ultrasound scanning via a plaster or catheter: Michiel Pertijs sees real possibilities for reducing the size of ultrasound equipment to millimetre format. And it's all thanks to smart chip technology.

*'Our collaborative partnerships enable us to make more progress towards a medical application. I find that extremely motivating.'*

Your very first photograph as a baby was probably an ultrasound image. A black and white picture on which, with a little imagination, you can recognise a baby in the womb. Ultrasound scanning is a cost-effective and safe imaging method for viewing inside the body. It is not only suitable for pictures of the unborn child, but for all kinds of organs, as long as they are made of soft tissue.

There are good reasons why doctors have long dreamt of probes that are so small that they can enable real-time imaging during minimally-invasive surgery, such as an angioplasty or the replacement of a heart valve. Or how about a wireless incontinence plaster that monitors bladder volume and sets off an alarm when it's time to urinate? All of this is achievable within the foreseeable future – in around 5 to 10 years, if you ask Michiel Pertijs.

## Miniaturisation

Together with its TU Delft partners, Michiel Pertijs' research group is working on the 'miniaturisation' of ultrasound scanning. This is because the technology needs to become considerably smaller before a doctor can fit a patient with an ultrasound catheter. Pertijs: 'For ultrasound scanning, you need a probe with thousands of transducers that transmit and receive ultrasound. Currently, this kind of probe is connected by a thick cable to a box next to the bed. This is packed with electronics: a piece of electronics to transmit and receive signals for every single element. However, passing thousands of cables through a catheter of 3 mm in diameter is simply impossible.'





*'Pertijs believes that he can make the chips even smarter. He is working on chips that digitise the ultrasound signals first, potentially making the flow of information even more efficient.'*

### Chip technology

The challenge facing Pertijs and his research group is to miniaturise the functionality of this large box of electronics and make the connection between the probe and the monitor as efficient as possible. Chip technology plays a role in this. 'We can use it to make smart switches at micrometre level,' explains Pertijs. 'Small enough to be able to control those thousands of transducers from the probe.' Thanks to smart switches, his chips can already process some of the signals before they are sent through the cables. 'If you can already combine the signals of groups of elements on the chip, you only need to transmit the sum of those signals rather than the signals from all those individual elements. This alone reduces the number of cables by a factor of 10.'

### Digitisation

Pertijs believes that he can make the chips even smarter. He is working on chips that digitise the ultrasound signals first, potentially making the flow of information even more efficient. 'You can compare it to the development of telephony: we used to make calls using analogue telephones. That telephone conversation is now digital. In the past, only this telephone signal passed through the same cable, but now we also receive broadband internet and around a hundred television channels directly into the home. Possibilities like that also apply in this field. Ultrasound probes are the old-fashioned analogue telephones. In medical equipment, the digitisation of data is still in its infancy.'

### About Michiel

**Michiel Pertijs** is associate professor in the Electronic Instrumentation Lab in TU Delft's Microelectronics department. He studied Electrical Engineering at TU Delft and obtained his doctorate in smart temperature sensors. He then became a product designer at the chip company National Semiconductor and a senior researcher at the Holst Centre in Eindhoven. After returning to research, he had the opportunity of a position as associate professor back at TU Delft. In it, he combines his expertise as a chip designer with ultrasound research.



## Wireless

All of that miniaturisation and reducing the number of cables: surely these chips could also communicate wirelessly? Pertijs: 'That depends on how much data you need to transmit. An incontinence plaster only needs to transmit an estimate of the bladder volume every ten minutes. That can easily be done wirelessly. But with real-time imaging, you're talking around tens of gigabits per second. That's a completely different order of magnitude.' In addition, wireless ultrasound can sometimes be counter-productive. 'A catheter always has a guide wire around it to push and pull it through the bloodstream. Some cables could easily run through a wire like that, making a wireless signal unnecessary. However, probes that the doctor holds by hand will certainly be wireless in the future. Our chips will be able to make an important contribution to that development.'

*'Ultrasound probes are the old-fashioned analogue telephones. In medical equipment, the digitisation of data is still in its infancy.'*

## Energy-efficient

But there is a downside to these increasingly smarter, smaller chips: processing all of the data streams on a mini-surface consumes a lot of energy. And this needs to go somewhere, in the form of heat. Although this is not a problem with large ultrasound equipment, it is not ideal in or on the patient. 'Heat can result in tissue damage,' says Pertijs. 'This is why we are working with our partners on energy-efficient, safe chips.'

## Cross-fertilisation

It is through collaboration with partners that Michiel Pertijs achieves the greatest progress. 'We are often working on the same problem from different perspectives. This means that we reach solutions that we would not have come to individually.' Using chip technology to create superfast ultrasound images is an example of this. This in turn can be used for new medical diagnostics. 'There is cross-fertilisation that you only find in this type of project. The bigger picture it gives us enables us to make more progress towards the medical application. I find that extremely motivating.'





## About ultrasound scanning

Ultrasound is known from the animal world: bats also use it. They emit ultrasound that reflects off their prey. That reflection – or echo – returns to the bat after a time. The time between transmitting and receiving is a measure of the distance between the bat and its prey. This principle can be translated to ultrasound scanning: a probe with a sound transmitter and receiver, a transducer, emits ultrasound. When the ultrasound signal hits a surface between two tissues, it reflects back. The transducer receives the echo and transmits it to the

ultrasound device, which combines several echoes into an image, for example of a foetus in the womb. An individual echo is comparable to a depth measurement from a boat: if you throw a lead weight on a rope into the water, you know how deep the water is by measuring the length of rope. By doing the same kind of measurement from several boats, you gain an image of the river or seabed. This is similar to a two-dimensional ultrasound image of an object in the body: the boat is the transducer and the rope is the route taken by the sound.





## A glimpse into the universe with superconducting electronics

# 29

..... Text: Dave Boomkens | Photography: Mark Prins

Questions like *How did the universe come into being?* and *Where do we, humans, come from?* can be approached in two ways: philosophically and scientifically. Scientifically this is done by looking at light – very old light. ‘And to do that you need superconducting detectors and sensitive measuring equipment,’ says Professor Jochem Baselmans. ‘Only then can you get answers to what may be the most relevant questions in our universe.’



*‘The Galileo example shows that new instruments lead to new discoveries, and that new discoveries lead to new insights.’*

Mankind has been trying to understand the universe for centuries. The early Babylonians made an attempt, but the best-known example is perhaps Galileo Galilei. In 1610, he used a telescope made by spectacle maker Hans Lipperhey from the Dutch town of Middelburg to gaze at the stars. ‘Galileo bought the telescope at a market in Venice, disassembled it and made it thirty times more powerful,’ says Baselmans. ‘Galileo spent days studying the four great moons of Jupiter until he was certain: the moons orbited Jupiter. His discovery demonstrated that not everything in space revolves around the Earth and undermined the geocentric world view that was considered true at that time. The Galileo example shows that new instruments lead to new discoveries, and that new discoveries lead to new insights. Since then, astronomy has been the branch of science in which new technology is applied the fastest by far.’

**New insights**

New instruments enable us to get a better understanding of the way the universe came into being. According to Baselmans, this technology is simply indispensable. ‘After all, you can’t just fly through the universe or study a star in a laboratory. Because observation is quite difficult in astronomy (huge distances, weak signals) and experimentation is not possible (having two galaxies collide to see what happens), technology is the key.’ Until the 1970s, the focus was on photographic plates, but development of Charge Coupled Devices, known as CCDs, marked the dawn of a new era: that of highly sensitive light detectors. ‘All of a sudden, astronomers were able to capture images and spectra of very weak objects. Just try to find a new planet or a new asteroid: you have to detect movement among one of those ten-million dots above you. This is much easier to do with a computer than with photos.’

**Best in the world**


This semiconductor technology can be purchased. ‘There’s a whole industry behind it,’ Baselmans explains. ‘It is used by Ministry of Defence, and also by commercial photographers. But much of the radiation from the universe is not visible light and is emitted at other wavelengths. An important example is submillimetre radiation, also known as far-infrared radiation. In order to detect this radiation, you need superconducting detectors and chips. It is not possible with semiconductor technology. But those superconducting detectors only work at very low temperatures; temperatures close to absolute zero (-273.15 degrees Celsius), which makes this technology too expensive and difficult for most applications. So you only use it if you really want to understand each ray of light, and that’s

**About Jochem**

**Jochem Baselmans** gave his inaugural address on 11 October 2019. He has a shared position, spending 50% of his time working for TU Delft and the other 50% for SRON. The combination of these two institutes makes it possible to develop revolutionary new instruments. The big challenge: developing integrated circuits based on superconducting electronics that combine radiation coupling, filtering and detection on a single chip. This makes it possible to create the broadband camera spectrograph mentioned in the text. This is also known as an imaging spectrometer.







.....

The Atacama Submillimeter Telescope Experiment (ASTE). A far infrared camera – developed in Delft – is installed inside the ASTE telescope.

*‘When I obtained my doctorate in 2002, telescopes like the Hubble – which now orbits the earth like a satellite – and ALMA were booming.’*

## Far infrared

Why is it so important to work with infrared light? What does far infrared tell us about the universe? ‘To answer those questions, I will first give you a short history lesson about the universe,’ Baselmans continues. ‘The universe was created about 13.8 billion years ago with the Big Bang – a hot ball of plasma that grew bigger and bigger in a very short period of time. With the exception of the afterglow of the Big Bang, the cosmic microwave background, the universe was still dark. Then the universe cooled down. The cooling of the universe enabled the formation of the element hydrogen, and the universe became transparent to light. With the birth of the very first stars, the lights went on in the universe. All the other elements were formed in the cores of those stars, and the supernova explosions that destroyed them. So we’re all stardust! With far-infrared light you can look at the oldest stars and all generations of stars that came into existence immediately afterwards. By studying that, you gain knowledge about the oldest galaxies, and also about the birth of planets and other galaxies.’

## Hubble & Alma

The Hubble Telescope was used to take many optical and near-infrared photographs of the universe. And we now also have ALMA: the Atacama Large Millimeter/ submillimeter Array, an advanced telescope that can detect far-infrared radiation from some of the coldest objects in the universe. In order to do that, ALMA uses superconducting technology developed partly at TU Delft and partly at SRON. ‘ALMA measures far-infrared light with a very high spatial and spectral resolution. To put it into simpler terms: ALMA can take very sharp pictures (spatial) and measure the precise type of radiation on each ‘pixel’ based on the wavelength (spectral). These spectroscopic observations make it possible to create a new image of the universe.’ Scientists looking to find out more about the early beginnings of our universe need to measure infrared light which has taken between 2 and 10 billion years to reach earth. Sensitive instruments are required for this. Jochem Baselmans is working on superconductive and extremely sensitive measurement equipment that can speed up the current measurement process 100-fold. Once it all works, we will be able to create 3D maps of star systems, allowing us to look back in time and space.


SCAN!



what I want! I want to help answer questions like *Where do we come from?* and that can only be done with superconducting detectors for submillimetre radiation. To make these detectors, I have a great partnership with the Terahertz Sensing group at TU Delft. This group includes two of the best in the world in the field of quasi-optics and electromagnetism, namely Andrea Neto and Nuria Llombart.’

## Missing link

According to Baselmans, the knowledge Andrea and Nuria have is the missing link in his expertise. ‘A telescope consists of one or two large mirrors that send light into the cabin of the telescope. To get that radiation into my superconducting chips properly, I need the knowledge of the Terahertz Sensing group. For example, one of Nuria’s PhD students designed the optics for DESHIMA – the Deep Spectroscopic High-redshift Mapper. DESHIMA, a chip about the size of two euro coins, is used to measure the distances and ages of distant galaxies. This is done by measuring different shades of infrared light. In October 2017, under the watchful eye of my colleague Akira Endo, researchers fitted this chip to the ASTE telescope in Chili.’



.....

Photo of the DESHIMA chip. Cooled to 0.1 degrees Kelvin, this small chip measures a spectrum in far-infrared. Previously, this could only be done with systems of about half a meter in size.



'I want to help answer questions like **Where do we come from?** and that can only be done with superconducting detectors for submillimetre radiation.'

### Imaging spectrometer

DESHIMA is another type of spectrometer that can very well be used in combination with a telescope like Alma. Although DESHIMA measures radiation at a much lower spectral resolution, it does so at a much higher bandwidth. This enables DESHIMA to determine the distance to far-infrared galaxies and to gather knowledge about star formation in dust-enshrouded galaxies. 'The nice thing about DESHIMA is that we now have a spectrometer that is as small as a lens, and it works. For the time being, we are still working with one pixel, but one of the next steps will be placing a large number of lenses side by side. Each lens will then have its own spectrometer. This way we will increase the number of pixels from 1 to 25. Or perhaps to 100! This then becomes an imaging spectrometer.'

### Into space

If the imaging spectrometer meets all expectations, the universe beckons. And to go into space, you have to be at SRON: the Netherlands Institute for Space Research, which is where Baselmans spends about half of his time at work. 'We also have to find someone who is crazy enough to build a telescope that can be cooled down to 4 Kelvin,' laughs Baselmans. 'That's about 269 degrees (Celsius) below zero. So you have to launch a fridge into space, which, with the help of a number of solar panels, is able to keep the whole mirror at 4 Kelvin. Otherwise, the radiation from the telescope will blind your detectors. What we are doing now works up to 1 Terahertz. But if you go into space, you want it to work up to at least 10 Terahertz. There's still a lot of homework for us there.'

### Cooler

So, superconducting technology, which only works in extremely cold conditions, can help us obtain an advanced picture of the universe. 'This ridiculously low temperature was also something that we had

to take into account when designing DESHIMA. For example, we converted and tested a cooler. In addition, you need reading electronics. Those electronics aren't something you can just buy, and there isn't any fundamental research either... no... that's high-level engineering. You also have to be able to understand the physics of the chips. I learned a lot about that from Teun Klapwijk, the expert in the field of nanotechnology. Once you have a good understanding of a chip and are able to read it, you still have to make sure the light reaches the right place. Otherwise, of course, it's no use to you at all. With current quasi-optical techniques, this is not possible for the large bandwidth of the future imaging spectrometers. Fortunately, we are working hard on that in EEMCS. In fact, over the past two-and-a-half years, my colleague Nuria Llombart of the Terahertz Sensing group has been working on finding out how this can be done at all. Turns out it's very difficult.'


### The power of EEMCS and SRON

But working in Delft and at SRON offers more, says Baselmans. For example, there are two large clean-rooms in Delft where all kinds of things can be tried out. In addition, SRON provides a great deal of knowledge and experience in the field of space instrumentation. He also wants to do more in the field of single processing in the near future. 'Once you have the ideal instrument, how do you then make sure you can make astronomical observations with it? How do you ensure that you're not affected by noise in the detectors or particles drifting through the atmosphere? And what if you want to work with 50,000 detectors at the same time? How are you going to get your data processed quickly enough? How do you ensure that algorithms work as efficiently as possible? The latter is something I don't know much about at all. Fortunately, a lot of that expertise can also be found at TU Delft and I can always pop in to ask my colleagues this question.'



.....


In October 2017, Dutch and Japanese researchers mounted this special chip on the Japanese ASTE telescope in North Chile. The superconducting chip is developed by Delft University of Technology and SRON, Netherlands Institute for Space Research. The chip contains one antenna, 49 filters and 49 detectors. The antenna captures radiation of various wavelengths. The filters unravel the radiation in 49 tones of infrared. The 49 detectors measure the intensity of the radiation. When a detector picks up a signal, it can be seen as a peak in a graph.



### Higher purpose

The reason for Baselmans' interest in superconducting electronics is related to his PhD research. 'When I obtained my doctorate in 2002, telescopes like the Hubble – which now orbits the earth like a satellite – and ALMA were booming. At that time, Terahertz technology with superconductors was cutting edge, while my PhD research was quite fundamental. What I missed while working on my PhD was an interesting context. I'm more into creating. I love making new things, but they have to be useful; they have to serve a higher purpose. Questions such as *How did the universe come into being?* and *How does everything fit together?* serve that purpose. It doesn't get more relevant than that.'





# We\Visit: combatting loneliness with video-calling

30

Text: Dave Boomkens | Photograph: Frank Auperlé

Visiting a loved one who's critically ill: normally that's possible, but due to the coronavirus, visiting ICUs is no longer permitted. The result is a lonely fight behind closed doors. In an attempt to combat the loneliness, researchers at Delft University of Technology have come up with We\Visit: a tool that facilitates video-calling by appointment. And it was developed at lightning speed after the Reinier de Graaf hospital sounded the alarm.

*'Our collaborative partnerships enable us to make more progress towards a medical application. I find that extremely motivating.'*

It all began on a Friday afternoon, shortly after Dutch Prime Minister Rutte announced the 'intelligent lockdown'. A researcher the university's Faculty of Mechanical, Marine and Materials Engineering (3ME) forwarded an email to Elif Özcan Vieira. Elif, who works for the Faculty of Industrial Design Engineering (IDE) and whose research interests include alarm fatigue among medical personnel, opened the message and immediately felt the hairs of her arms stand on end. 'It was a pretty emotional email,' Elif explains. 'A cry for help. The Reinier de Graaf hospital said that it was in dire need of a communication tool for its ICU. The patients were not only suffering from the coronavirus, it said, but also from loneliness. That's a heartrending combination, and I immediately felt that I had to get to work. The research experience I'd gained at various ICUs would certainly be useful, but I needed help; I needed technical knowledge to be able to develop such a tool.'

## Interfaculty cross-fertilisation

A colleague put Elif in touch with Willem-Paul Brinkman. Working just 650 metres along the campus, Willem-Paul uses virtual technologies in his work combatting phobias and psychotic disorders at the Faculty of Electrical Engineering, Mathematics and Computer Science (EEMCS). That very same weekend, he drummed up a number of fellow scientists, who attended a Zoom meeting with him and Elif the following Monday. One of them was Merijn Bruijnes, a psychologist specialising in the interaction between people and technology. Merijn: 'The email from the Reinier de Graaf hospital created a stir, and immediately led to a brainstorming session. We soon came up with the idea of video-calling including a planning tool. But the question remained: who would we be making such a tool for?'





### No mobile phones

In order to reach the right target group, the researchers contacted, among others, the FCIC – a Dutch foundation that promotes the interests of ICU patients and their loved ones. Elif: ‘Often people admitted to intensive care don’t have access to a mobile phone. After all, they had to be rushed to hospital. And there’s also the fact that corona patients are often aged 65 or older.’ Merijn adds that many corona patients in the ICU are in an induced coma. ‘Some of the patients have been intubated and are attached to breathing apparatus. That looks very intimidating, but someone from the foundation explained that this doesn’t necessarily mean that people are no longer taking anything in. So how do you adapt your tool to that?’

### Cybersecurity

The brainstorming session produced a very specific concept: in order to create a feeling of social contact and provide emotional support to loved ones, the scientists wanted to make an online video tool. On a number of points, this video tool would take up where FaceTime and WhatsApp left off. Not only can the tool be used easily by those aged over 65, for example, but protecting patients’ privacy is also an important element. ‘The Cybersecurity group put forward a number of students who did their best to hack the tool’s security protection. They reported their findings and then came up with solutions. That proved to be really helpful.’

### Prototype

One week after the Monday afternoon session, the researchers had a prototype. Its name? We\Visit. And how does it work? Merijn explains: ‘The family’s contact person is given a code by the nursing staff. The patient’s family

‘Protecting patients’ privacy is also an important element.’

### About Elif

**Elif Özcan-Vieira** is the director of TU Delft’s Critical Alarms Lab (CAL), and is the Care Technology Lead at the Adult Intensive Care Department of Erasmus Medical Centre. In the article *A beautiful alarm beside your hospital bed*, Elif explains that a cacophony of sounds can cause alarm fatigue among medical personnel and anxiety in patients. In collaboration with the Erasmus MC, she aims to address the problem of noise pollution in ICUs with human-centered design solutions that rely on advanced technologies. During her student days in Ankara, when she studied Industrial Design, she fell in love with both radio and sound. ‘There’s a mysterious magic to sound: its shape, its apparent elusiveness, its direct effect on the immediate surroundings. As a sound designer at Radyo ODTÜ, which was then being run by students and has now grown to be a national station, I became fascinated by the tiny vibrations that we catch with our ears. In the years that followed, I began to learn more about product perception, human behaviour and psychological processes – and sound kept cropping up in that context, too.’







can use this code to log in to the We\Visit system to schedule a virtual appointment. When it's time for the appointment, someone from the nursing staff – or a volunteer – takes a tablet to the patient to start the video connection at their bedside. As with Zoom, there is a virtual meeting space where people can talk to each another. The tricky thing is that you're dealing with an ICU environment, which means that many things are not possible. In the beginning, we were still thinking really big, but we eventually came up with quite a simple design, developed by IDE student Marek Torbus and implemented by EEMCS students, all on a pro bono basis and in addition to their exams. In no time at all, my colleague Otto Visser (of the Distributed Systems Group) had ten students in his app saying they really wanted to work on this. It was incredibly special.'

#### About Merijn

**Merijn Bruijnes** is a psychologist and member of the Interactive Intelligence group at TU Delft. 'When it comes to computer science, people often think of hardcore programming or machine learning, not psychology. It's nevertheless the case that you can use psychological knowledge to make technology work more effectively. Take the example of making a system user-friendly. More and more organisations are working with chatbots or virtual assistants. By researching how people respond to digital interlocutors such as these, technology can be used to contribute to psychology.' One good example of how psychology and technology can come together is the tool that Merijn developed for the Police Academy in Apeldoorn. He devised a virtual suspect that could be used by trainee police officers to practise various cross-examination techniques. In this way, there is no need for up-and-coming officers to exchange their textbooks for the interrogation room straight away; instead, this provides a great intermediary step. Merijn's current research focuses on diabetes patients. 'Often, we only consider the patient and the doctor, but that means we miss the bigger picture. That's because patients also interact with a social environment. 'How can I tell my daughter that she's allowed to bake cookies, but I'm not allowed to eat them?' I'm therefore mainly looking at how you can improve this exchange between patients and their social environment.'

*'What do healthcare staff require? And what do family members at home want?'*

#### Open science

The application used for video-calling, Jitsi, is a so-called *open source application*, which means that anyone can use it. And sharing is a method that fits perfectly with the We\Visit team's approach. Elif: 'At the moment, we're seeing physical borders closing all over the world, and countries rapidly turning inwards to help themselves. With this open science approach, we want to transcend national borders. It isn't just our virus, but everyone's virus. Although we're making a technical product, this is all about the need to solve a social problem. That's why open science is so incredibly important during this corona crisis.' This vision is endorsed by Merijn: 'From the outset, we all agreed that everyone should be able to use our findings. At present, we're running a We\Visit pilot at the Reinier de Graaf hospital, and we'll soon be able to test a more finely tuned version at Rotterdam's Erasmus Medical Centre. If we're satisfied after the pilots, we want to release the software, so that any stakeholder from any organisation will only need to download a package of information. A ready-made package will soon be available on our website.'

#### Packages

Elif has first-hand evidence of the need for ready-made packages such as this. 'Last week I spoke to one of my uncles, who is developing tests to detect hepatitis in a lab in Turkey. We discussed the current pandemic that's setting the world ablaze. When I told him about We\Visit and the software packages that everyone would soon be able to use, he immediately said: 'OK, let me know when it's ready, because this has to come to Turkey.' We hope that soon, We\Visit will not only be used in hospitals, but also in care homes and nursing homes, both in the Netherlands and much further afield.'





### Bringing a smile

Until then, however, there's a lot of hard work to come. Elif: 'We've received some funding from the TU Delft COVID-19 Response Fund, but we'll soon need to think about long-term financing too. After all, we could do so much more with We\Visit. We could create a 'message in the bottle' function, for example. We see that young children in particular can become extremely shy when faced with a camera. They would much rather sing a song for granny in a familiar context, or take the time to create a cheerful drawing. With a 'message in the bottle' function, the family wouldn't need to come together at the same time, but you could send different messages throughout the day. We\Visit would combine them and deliver them as a 'message in a bottle'. But developing something like that will take time.' Merijn: 'In order to put all these great ideas into practice, we will indeed need a business plan. Our colleagues at the Faculty of Technology, Policy and Management (TPM) might be able help us with that in due course. I do believe that it is going to work. When you see how willing everyone is to get involved, it's really inspiring. For now, the key thing is to reduce the pressure on the ICUs. Doctors and nurses are currently working one exhausting shift after another, and until now there's been little room for the human aspect. We hope that We\Visit will change that, and provide a reason for a smile without having to sacrifice anything – or anyone.'

*'We see that young children in particular can become extremely shy when faced with a camera. They would much rather sing a song for granny in a familiar context, or take the time to create a cheerful drawing.'*

The design of We\Visit, developed by IDE student Marek Torbus, is deliberately simple.

### About Cédric and Marek

**Cédric Willekens** is studying computer science at TU Delft, and was introduced to the We\Visit project by Otto Visser. 'Shortly after COVID-19 took hold in the Netherlands, I moved in with my girlfriend. We often talked about her grandmothers, and how to keep in touch with loved ones at a time of social distancing. When I heard Otto Visser talking about We\Visit, I immediately knew that it was a fantastic project and that I wanted to contribute. The great advantage is that, as a student assistant, I've been working on educational software projects for some time now. One of those projects is called Queue. Queue is a tool that you can use to schedule contact time with students; if there's an oral exam coming up, for example. The basic idea of Queue is similar to that of We\Visit. At the outset, in very close cooperation with the Reinier de Graaf hospital, we tried to make an inventory of the needs. What do healthcare staff require? And what do family members at home want? Based on this input, we wrote the back-end – the part that's not visible to the user. The industrial designers focus on the user interface and try to make the digital environment as user-friendly as possible. That includes asking questions like: what buttons does the device need? Which route will the user take? They make things intuitive, make sure there's a sleek design, and form the link between us – the computer scientists who are responsible for the technical implementation – and users.'

The design of We\Visit, developed by IDE student **Marek Torbus**, is deliberately simple. Marek: 'I was really interested in doing this project because of its current need. I talked to Elif about doing something to help and was lucky enough to join the team from the beginning. My main responsibility was to create a workflow and user interface for our service. The main requirements were to make the website accessible for different users, especially taking into account the elderly who have difficulties in dealing with technology or don't have a smartphone which makes it very difficult for them to communicate with their family and friends when in the ICU. In order to fulfil this, the website has been designed with the minimal possible steps needed to connect families with the patients. And I also used for example very contrasting colours. The key to our success has been daily Zoom meetings where I had the opportunity to present work and receive immediate feedback from the team. I'm particularly honoured to be a part of We\Visit.'





# EEMCS NODES



## VOLUME II

Copyright © 2020 TU Delft

**CONCEPT AND DESIGN**      Hollands Lof, Haarlem

**PRINT**                              Veenman+, Rotterdam

NODES is a publication of the Faculty of Electrical Engineering, Mathematics and Computer Science.  
No part of this publication may be reproduced and/or made public without written permission.







# EEMCS NODES

## VOLUME II

This book contains a time-based selection of the scientific research that takes place at the Faculty of Electrical Engineering, Mathematics and Computer Science (EEMCS). EEMCS is the faculty where bits and bytes are converted into real solutions for people and society. And that is desperately needed. Because our safety, health and well-being depend more and more on technology. The scientists of EEMCS are working, day-to-day, on smart solutions to global problems and do not shy away from any challenge. They change the world and make impact. They are our heroes. They tell our stories. They are the creators of the future.

 [tudelft.nl/eemcs](https://tudelft.nl/eemcs)

 [@EEMCS\\_TUD](https://twitter.com/EEMCS_TUD)

 [eemcs\\_tud](https://www.instagram.com/eemcs_tud)