MAXWELL

Edition 26.1

FLOW

Research
Data transport in radio astronomy
Eduction
New bachelor curriculum
Association
Stories about EEMCS

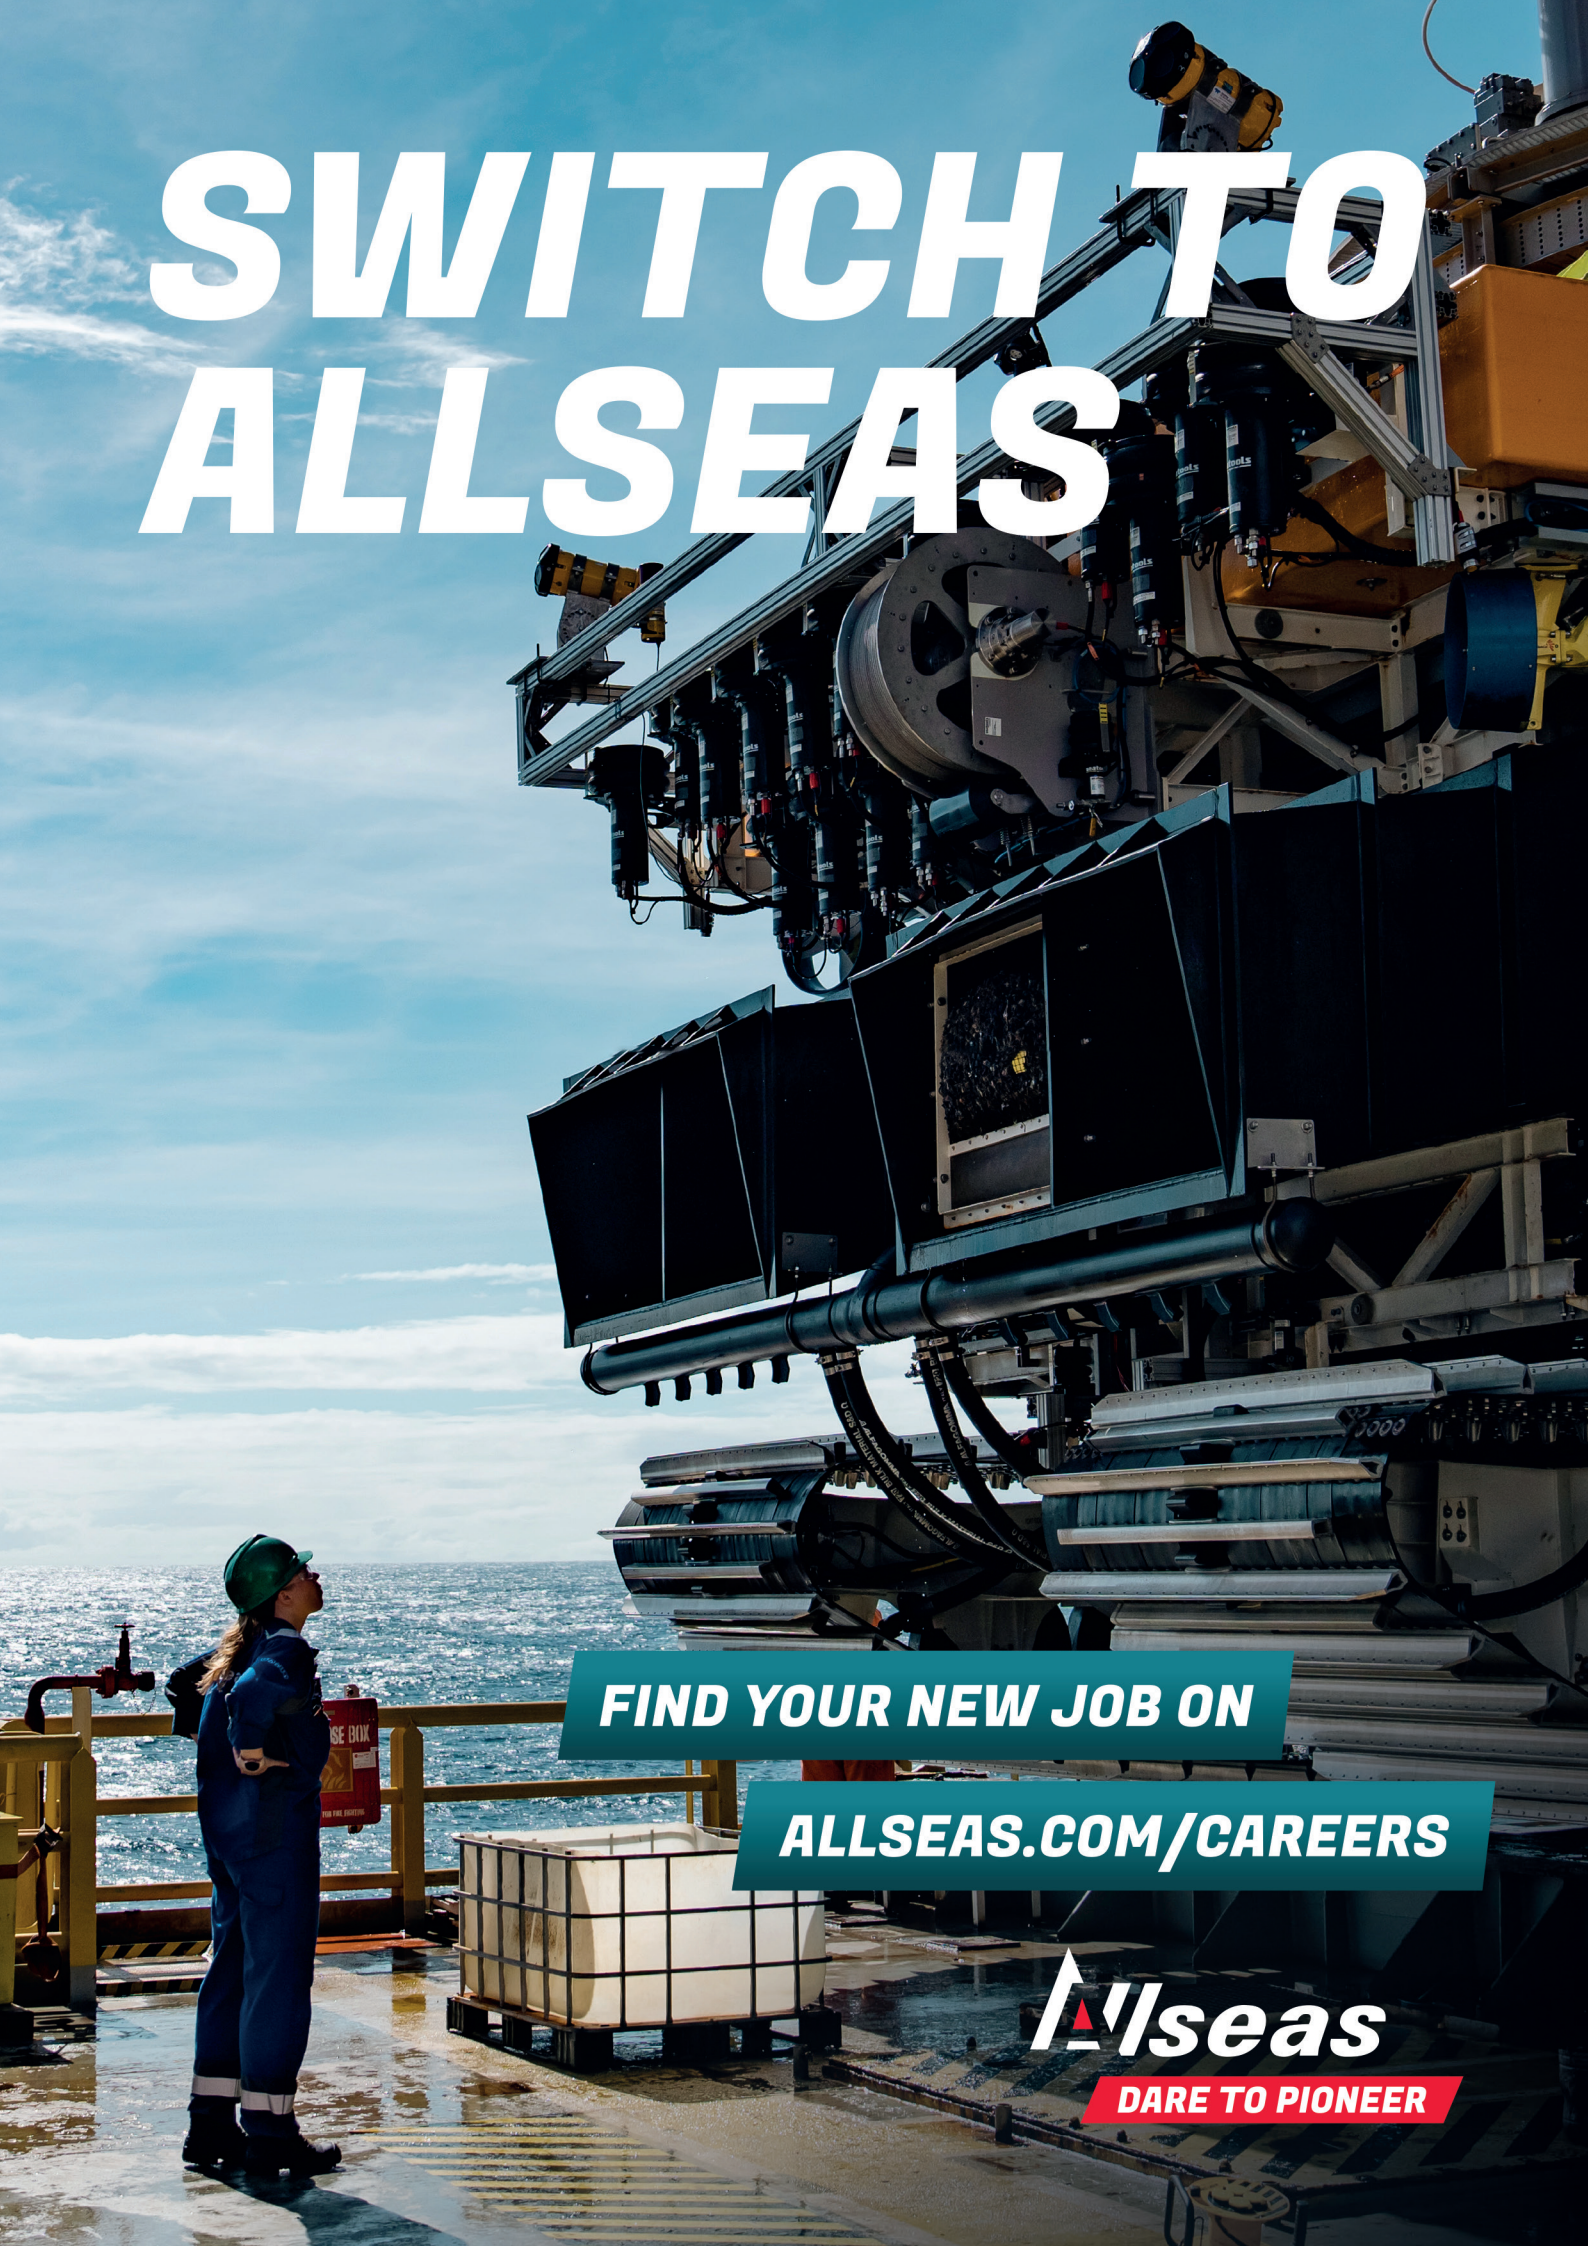
SWITCH TO
ALLSEAS
FIND YOUR NEW JOB ON
ALLSEAS.COM/CAREERS
Allseas
DARE TO PIONEER

From the Board

Dear reader,

The past academic year the 150th Board has brought new light to the ETV past COVID times. They had introduced a wide scala of activities, with the lustrumweeks and gala, and with new initiatives such as the two day rally weekend and the batavierenrace participation (where I myself got up at 5 in the morning to run 10 km). Furthermore they introduced a new Board member, the commissioner of career affairs. Nevertheless, they kept the small traditions and gestures that came forth from our rich 116 year long history and organised all the activities that have long been traditional in the ETV.
This is how the 150th Board set a great example of the way that tradition and innovation can seem like a paradox, but are actually two interwoven concepts. Without a fundament there is no room for improvement and innovation, but without improvement the fundament does not remain solid.
We, the 151st Board have also begun this year with that perspective in mind. The past half year we went ahead and had a Fresh look at the ETV. The core values of our association of course remain: education, career and social connection between the students. However, we have been striving to make the ETV a more diverse place. To do so, we have been actively including master students more and we've promoted our events more directly to them. We found it a great idea to pay extra attention to sports (and alcoholfree activities) and so we've re-introduced the sports committee, who had a great run with their dodgeball activity and is going skiing with our members soon. Furthermore, we put our efforts to creating a more conscientious association, especially on the topics of sustainability and inclusivity, which we are realising through putting the topic of sustainability on every committee's standard agenda and by introducing (as first study association in Delft) the inclusivity committee, who are promoting the interests of the LGBTQ+ community.
So far, we are holding on to the core values and traditions of the ETV while implementing them with a Fresh look. The beautiful thing about our Boardyear is that we really feel like everyone is granting us the freedom and support to do so and to innovate in the direction we want. We are confident that together with our members, and former Boards we will make the rest of the year unforgettable and amazing. We wish you happy reading of this Maxwell still!

With Fresh regards,

On behalf of the 151st Board of the Electrotechnische Vereeniging,

Simon Molenkamp

Colofon

Year 26, edition 1, January 2023
Maxwell-ETV@tudelft.nl

Editors
Anne Hinrichs
Rui Mao
Savion Dmello
Joris van Breukelen
Thomas Pouwels
Joos Vrijdag

Contributers
D. Muratore, J. Moeder, N. aderintevullen

Advertisement
page 2, 18,19 - Allseas
page 24,25 - ASML

Printing
Quantes, Rijswijk, 1400 copies

Electrotechnische Vereeniging
Mekelweg 4
2628CD Delft
The Netherlands
+31 15 278 6189
board-etv@tudelft.nl
www.etv.tudelft.nl

151st Board of the ETV
President - Simon Molenkamp
Secretary - Joran Kroese
Treasurer - Bart Zuidema
External Affairs - Maya Peeters
Education - Julie Diender
Carreer - Mirthe Otter

Subscription
non-members can receive the Maxwell two times a year, against a contribution of €5,- per year.

Change of address
Please send your changes to the (email) address above or use the website.

Editorial

Dear Reader,

Welcome to the first edition of the Maxwell's 26th volume! With this edition, we have decided to revamp the format. By having dedicated association and education sections, we aim to make the Maxwell more attractive to a broader audience. Another change we felt was necessary was to limit ourselves to only two editions this year, allowing us to produce higher-quality articles and a denser Maxwell. To signify this gradual change we have chosen this edition's theme: Flow.

The magazine in front of you contains the latest news in the Electrical Engineering field, with articles on energy harvesting and FPGA programming. It also contains a short history of post-WWII Dutch encryption technology, provided by the Historic Collection, with exciting discoveries and a surprising link to one of the Netherlands' first computers. Next to that, we have thought-provoking interviews on the new bachelor program and the CESE master, and an interview with Marc Vlek, who wrote a book on his family history which is also the history of the ETV's favourite drink. After discovering why EWI used to be Delft's best base jumping spot, we reach the highlight of this edition: the Maxwell puzzle section, containing the ETV's first (after) Christmas Puzzle which has a prize worth €230!

On behalf of the Maxwell Committee, I would like to thank all guest authors and contributors, without whom this edition would not have been possible. I would also like to thank my fellow Maxwell Committee members and the ETV Board for their dedication and drive to deliver this new Maxwell format. Lastly, sadder news. With the passing of ir. Rob Timmermans, the Historic Collection loses a valuable member. We express our sincere condolances as he was an important contributor to previous Maxwells.

We wish you a pleasant reading experience!

Anne Hinrichs

Contents

Research

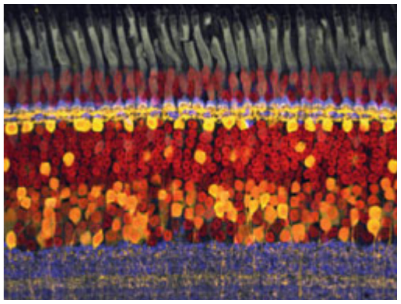

6 The Stanford Artificial Retina

10 Data Transport in Radio Astronomy

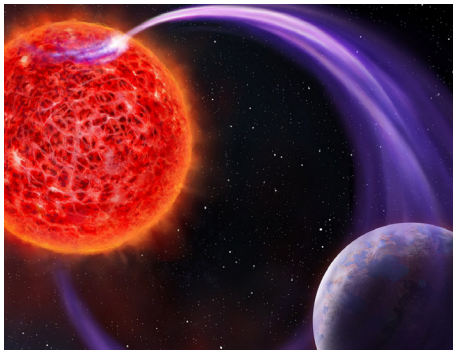

17 Energy Harvest Circuit

19 Study Collection

Education

26 New Bachelor Programme

28 New Master CESE

Association

30 Stories of EEMCS

32 Interview with Marc Vlek

34 Capture the flag

35 Puzzle with a price

36 Flowchart

37 /Pub Beer Review

38 Activities

The Stanford Artificial Retina

An interview with dr. D.G. Muratore

By J. Vrijdag

One of the most exciting topics in electric engineering is the brain computer interface (BCI), which allows the digital domain and the biological brain to connect. There are many types of these BCI's, one of the most futuristic among these is the artificial retina. The general idea behind such a device is to replace the natural eye with a camera, which allows vision for certain forms of blindness. We sat down with a researcher from TU Delft who is doing research in this field, dr. Dante Muratore. His research is in collaboration with Stanford university to make a new type of artificial retina.

So first of all, thank you for being here today. Can you explain to us what an artificial retina is?

The goal of an artificial retina is to replace the natural functionality of neurons in the retina for people who suffer from profound blindness. The retina (Figure 1) is a multi layer structure that sits at the back of the eye and is responsible for transducing the incoming lights into spikes. These spikes form patterns of neural activity that are then transmitted to the brain. Therefore we will think of the retina not only as a camera that produces an activity that is proportional to the incoming light, it also has another layer. That first layer is called the layer of photo receptors that, oddly enough, sits at the back of the retina. The light goes through the retina, hits the photoreceptors, which then transduce that light into electrical activity. Next there are a couple of layers of neurons that do some local processing before the signal gets send to the brain. This local processing is done by interneurons, which do processing in the analog domain, and retinal ganglion cells (RGCs) to do processing in the digital domain, meaning that they are spiking neurons. So each of these neurons generates an action potential, a spike. Now this is interesting, because those action potentials can be recorded and you can learn a lot about the retina. What we already know is that there is not only one type of neuron, there is actually multiple. Those different types of neurons contain different types of information. You have, for example, neurons that can differentiate between on and off. The on-type neurons will start spiking when there is an increase in light and the off-type starts peaking when there is a decrease in light.

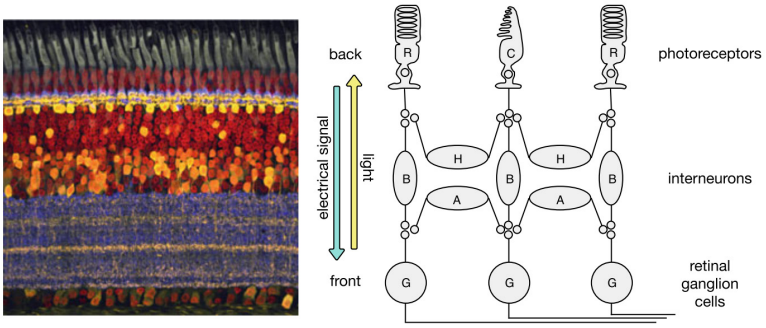

Figure 1. Left: retina cross section. Right: schematic of retina layout and cell classes: (R)od and (C)one photorecepters; (H)orizontal, (B)ipolar and (A)macrine interneurons; (G)anglion cells (from [1], © Abrams 2010).

There are two main types of artificial retina, one is called epi-retinal (figure 2), This is what we are working on, and the other is a sub-retinal. The difference is basically in the position of these implant devices. They both are artificial devices that we implant in the eye. The epiretinal system is on top of the layer that is the ganglion cells, the spiking neurons, while the subretinal is below the retina and it sits close to

> *"We expect that that will create a representation of the external world in the patient, who is typically blind"*

the photo receptors. In the subretinal case, basically what you're trying to implant is very efficient photodiodes that do the same as the photoreceptors. These photoreceptors are incredibly efficient. You can imagine this by realizing that you, a person without any vision impairment, can see very well while staying outside on a sunny day, which is not very common in the netherlands. On the other hand you can also see very well in the dark. So the dynamic range of the photoreceptors is huge, right? We cannot mimic that very easily with electronics, but there are people playing certain tricks to try to improve on that. That is one approach and our approach is called

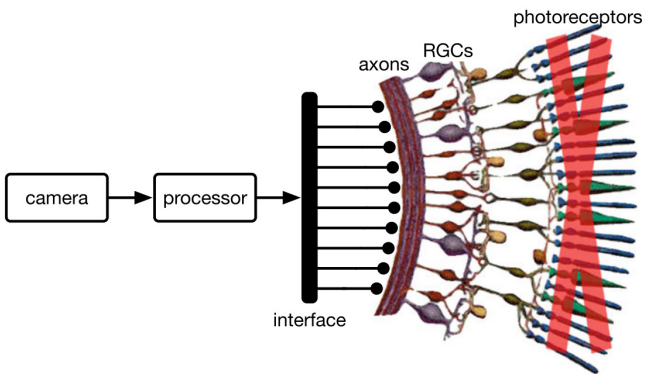

Figure 2. Schematic of retina with degenerated photoreceptors and epi-retinal prosthesis (from [2], 2020).

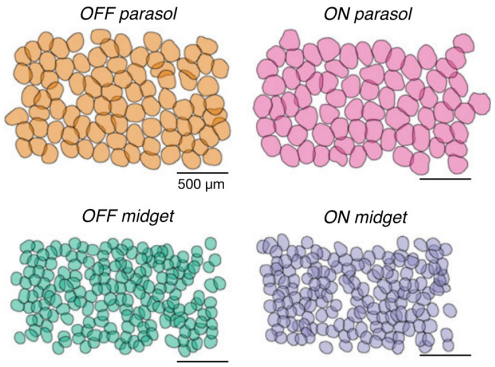

Figure 3. Receptive fields of four major RGC types (from [3], © Cell Press 2019).

epiretinal. This comes from the other side and it directly elicits the ganglion cells, the spiking neurons. So we try to stimulate them and the neurons generate a spike and the spike then travels through the optic nerve to the brain and is further processed. We expect that that will create a representation of the external world in the patient, who is typically blind.

And the subretina are on- and off-cells?

so basically they replace the photoreceptors and activate the bipolar cells and the whole part of interneurons directly. They don't directly stimulate the ganglion cells and most importantly, they cannot record from the ganglion cells. That is a key differentiation in what we do. Typically when people think about artificial, like retinal, processes is that you want to stimulate neurons. And that is true, but at the same time, the argument we make is that you want to know what you're stimulating. This goes back to different cell types (figure 3), because if you activate at the same time in the same location both on- and off-cells without realizing that you're telling the brain that that was the same time and the same place, that is a conflicting message. You don't want to do that. Otherwise the messages you are sending to the brain are very confusing. Our goal is to activate not only with single cell resolution, but also with cell type resolution. We want to know that we are activating this neuron here that is a certain type and not the neuron that is ten micron to the left or ten micron to the right that is a different type. For that we need a bidirectional interface (figure 4).

That is the stanford model, you're referring to?

Yeah, so this project was initiated by E.J. Chichilnisky, who is a professor at stanford university. That is where I got involved with it, back when I was a postdoc at stanford. I started developing some of the implantable electronics for that project. Now my lab is here, it's three phd students and four master students working from here together on the project. Since I moved here, we are an external collaborator, but we still work together basically.

So the main difference between the stanford model and the other devices is the measuring and the dictionary?

Exactly, so our goal is to build what we call a dictionary. but think of it as a matrix that tells you if you pass this much current on this specific electrode, what the probability is of activating any of the cells. Ideally you would have one hundred percent probability for the cell you want to activate and zero for the others, which we can achieve very often. Nevertheless there are some cases where maybe you activate two cells or more. In that case there are all sorts of algorithms to figure out based on the image, what is the best combination of electrical patterns of stimulation you want to pass, so that the image will be approximated the best. To do that, we need the dictionary.

To build the dictionary, we need a bi-directional interface that can stimulate and record at the same time. This is a technological challenge because, you stimulate and record on the same electrode. Now if you think about the electrode, one thing to notice is that the electrode has a very high impedence. This means that when you pass even a tiny amount of current on an electrode, you generate a large voltage on the same node where you want to record your tiny spike. This means you are kind of self interfering with your own recording. This voltage you're generating is what we call a stimulation artifact. But we want to see a very tiny spike and we cannot not generate the spike without the artifact, because otherwise we don't stimulate. This is wy we want it to be gone very fast, so one of my master's students recently developed an algorithm. Now a phd student is working on further moving this research into how to shape the stimulation to reduce that artifact as fast as possible, so that we can see the spike coming along.

This is one of the hundred problems we are targeting in parallel, which makes this very exciting. It's typically hard to find a project that has such high impact on society from a

Û

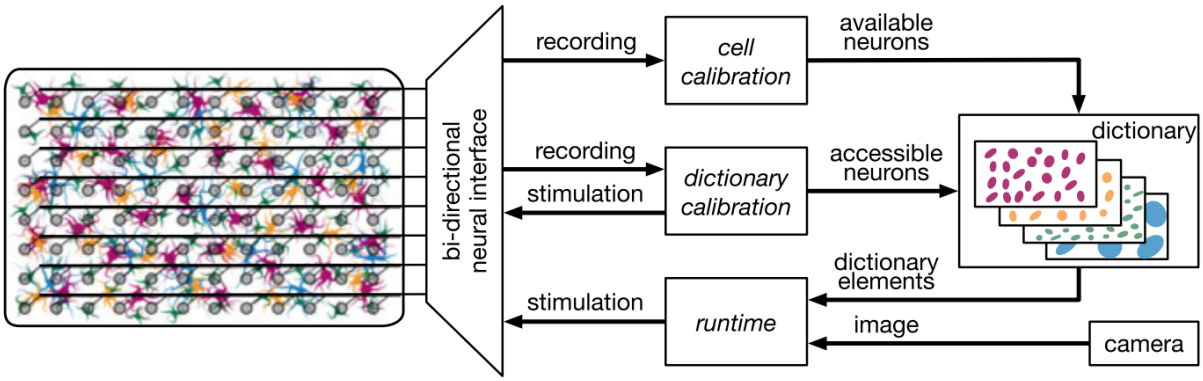

Figure 4. System level of the proposed bi-directional neural interface with three operating modes (from [2], 2020).

medical point of view that also has many engineering challenges that are worthy of research. Very often, the technology that is needed is quite simple, so it's not interesting. You cannot justify as an engineer who's paid for doing research to spend time on that, or at least a significant amount of time. Though in this case it is very nice, because the project is very fascinating and the challenges that come for electrical engineering are also very interesting.

Is there anything else about the stanford model that you're working on, that you would still like to share?

What makes this whole project very promising, is that we want to build an implant, that already exists in a lab for in vitro experiments. The lab routinely does this experiments with ex vivo tissue. So everything we want to try, we can try in that setup before we have to go into people. This is a great relatively quick turnaround for the technology, so we can design the electronics for the implantable device much better this way. It's very rare to have that. I think that is the main advantage that we have.

How many years do you think it will take before this can replace the human eye?

This is always a hard question, I don't know the answer, but there are various steps. I think that my goal is to be viable in animal experiments in five years. Meaning that after these five years you can do a lot with animal experience in terms of proving safety, but also efficacy. You can for instance train a monkey, to recognize a face in a picture. You train this monkey with a working eye and then in the other eye there is the implant. You try to see how well do they do the same task with the implant. Then you can stimulate in the clever way that we want to stimulate in or you can stimulate in, let's call it a traditional way, with one electronic to one pixel. What we want to prove is that respecting the natural patterns of activation of the retina, so this cell type pacification, will improve the performance a lot. You still cannot directly measure how well each solution works. We can however measure that we are activating that neuron, but the question is, and this is a huge hypothesis, that it matters. Until we build this and we show that it matters we cannot prove it. In order to prove it, we will need to have a model of how the brain decodes, and that is a complicated matter. Proving that it works is kind of like a chicken and egg problem. the only way to prove this is to build the technology. We do however have strong evidence that it makes sense. But the ultimate proof will be to stimulate in the way we propose and stimulate in the traditional way and show that our way allows to reconstruct much better image quality. I mean, we can do this in animals first, but the final proof is the human. In five years we will get to that animal experiment, that's my goal at least. From there it is unknown because it is not only about the technology development, but it's also regulatory problems. I am no lawyer, so I have no idea, but it it could take another ten years before we test in with humans or maybe five. It's also a matter of how much money you put into it. At that point, I think that this is going to be a project that doesn't belong to the university anymore, but it should be carried over by a company. Once we prove the science is solid and it just becomes regulatory and development, then it should be carried out by a company.

Ê

[1] C.E. Schoonover,Portraits of the Mind: Visualizing the Brain from Antiquity to the 21st Century. Abrams (2010)
[2] Muratore, D.G., Chichilnisky, E.J. (2020). Artificial Retina: A Future Cellular-Resolution Brain-Machine Interface. In: Murmann, B., Hoefflinger, B. (eds) NANO-CHIPS 2030. The Frontiers Collection. Springer, Cham. https://doi.org/10.1007/978-3-030-18338-7_24
[3] C.E. Rhoades, N.P. Shah, M.B. Manookin, N. Brackbill, A. Kling, G. Goetz, A. Sher, A.M. Litke, E.J. Chichilnisky, Unusual physiological properties of smooth monostratified ganglio cell types in primate retina. Neuron 103, 658–672.e6 (2019)

A deeper dive: the challenges of spike measurements

Next to the interview, we would like to provide a closer look to one of the challenges regarding the artificial retina project. Thanks to a chapter written by D.G. Muratore and E.J. Chichilnisky [1], some more exiting technical details about RGC spike measurements can be found below. The following text is a part of this chapter, edited by the Maxwell redaction.

In order to get a sufficiently large data set to be able to classify the electrical features of the different RGC's, many channels that record are needed. Around 10,000 channels with 10 bit resolution at 20,000 samples per second generate around 2 Gbps of recorded data. This can however be drastically reduced by exploiting the fact that the spikes that need to be measured are sparse, both temporally and spatially. The recording proposed in the Stanford artificial retina performs lossy compression before full digitization. This compression uses not only the fact that the spikes are temporally sparse (and therefore does not sample in between spikes), but also that only the distinction between spikes from different cells matter. This distinction can be made without a perfect representation of the spikes. Recording these properties can be accomplished efficiently using a *ramp-analog-to-digital converter (ADC)* coupled with a *wired-OR readout* and a *unique-signal decoder*. For each sample the ramp ADC indicates the measured voltage with a short pulse at a discrete time step proportional to the quantized voltage (The number of discrete time steps sets the ADC resolution and the ramp voltage range sets the full-scale range). This is achieved by comparing the input signal to a ramp voltage that steps through the entire input range (figure 1a). The time of the pulse is captured using a counter that keeps track of the ramp steps. The ramp and counter can be shared between all channels, making this an efficient algorithm for digitalizing many channels in an array. Then the channels are combined with an OR logic, across the rows and across the columns (figure 1b), to achieve the desired compression. Consequently, if only a single channel produces a pulse at a given time step (i.e., it is the only channel with a quantized voltage corresponding to the time step), then the the channel location is indicated by a uniquely decoded row and column. When multiple pulses occur at once, multiple channels and rows are activated and no uniquely decoded channel is indicated. Only uniquely decoded channels are stored, leading to substantial compression.

Ê

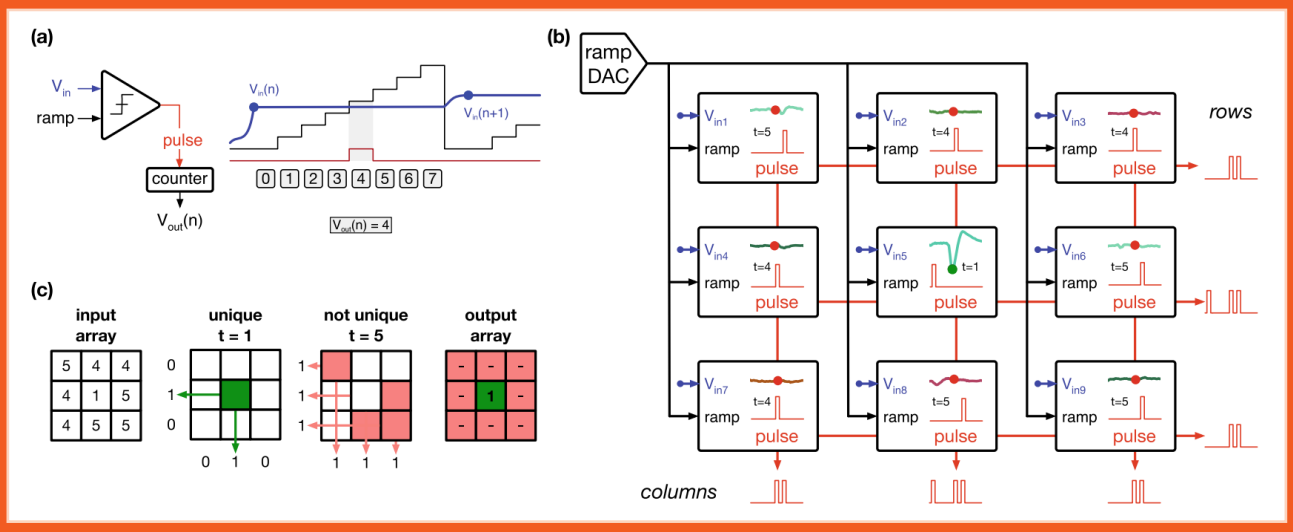

Figure 1. Compressive readout strategy. **a** Ramp ADC: schematic and operation. **b** Wired-OR readout: 3 × 3 array example. **c** Unique-signal decoder: only the channel(s) presenting a unique digitized voltage is recorded in the output. Examples of unique and non-unique activated locations at two time steps are also shown (from [1], 2020).

[1] Muratore, D.G., Chichilnisky, E.J. (2020). Artificial Retina: A Future Cellular-Resolution Brain-Machine Interface. In: Murmann, B., Hoefflinger, B. (eds) NANO-CHIPS 2030. The Frontiers Collection. Springer, Cham. https://doi.org/10.1007/978-3-030-18338-7_24

Data Transport in Radio Astronomy

Exploring the application of Remote Direct Memory Access over Converged Ethernet (RoCE)

By ir. W. de Laat, ir. S. van der Vlugt (ASTRON), dr. ir. Z. Al-Ars

In the past century, many discoveries have been made in the universe through radio astronomy. Several discoveries of recent years include Pulsars, Fast Radio Bursts and the magnetic interaction between stars and their planets[1] as illustrated in Figure 1. Contrary to an optical telescope, a radio telescope observes a range of radio signals outside the optical spectrum. Today's radio telescope systems consist of several to many antennas which are computationally combined into one large radio telescope. To enable more groundbreaking research, even better systems are needed to reveal more details of the universe.
At the same time, there is a pressing need to improve the material cost and energy efficiency of systems. One of the biggest challenges we face today is efficient data movement from antenna receivers to the GPUs in the central processor that performs most of the computational work. In this article, we discuss how an implementation with Remote Direct Memory Access (RDMA) over Converged Ethernet (RoCE) can both improve the energy efficiency of radio telescope systems as well as enabling new science.

Radio telescope systems

To enable groundbreaking research in astronomy, cutting-edge observation instruments need to be constantly improved. Integrating new technologies in such instruments must meet several tight constraints. The physical location of an instrument, for example, sets limitations on its size and energy budgets.

At the same time, the environment imposes conditions on the required robustness due to the impact of weather conditions and wildlife. In addition, these are often one-off large system designs expected to last for decades, making maintainability and reliability even more important. These requirements must be fulfilled within tight financial budgets for development, production and maintenance.

Alongside these practical tradeoffs for production systems, there is a tradeoff between the physical and computational properties. To go beyond the currently available high sensitivity and high resolution, one would need a very large (100 square meters to several square kilometres) collecting area. These kinds of systems are very costly and impractical to build. An alternative that represents the current trend is to build interferometer (i.e. phased array) systems consisting of many small antennas spread out over a large area. The signals of these antennas are computationally combined into one large antenna, shifting the physical and mechanical problem to a digital data processing problem.

An example of one such system is ASTRON's LOFAR [1] (LowFrequency Array) which is the largest low-frequency radio telescope worldwide. LOFAR consists of many receiver stations spread out over the Netherlands and Europe, as shown in Figure 2. A toplevel view of the LOFAR system is shown in Figure 3. These stations stream high volumes of data over pri-

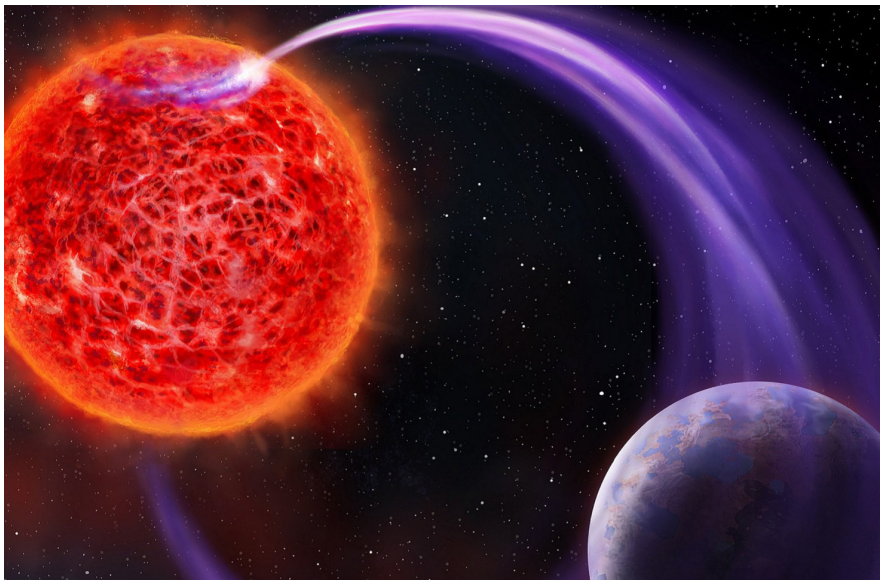

Figure 1. An artist's impression of star-planet interaction. The plasma flowing between the two bodies is depicted in violet. CC-BY-NC-SA (Credit: Danielle Futselaar).

[1] https://www.astron.nl/dailyimage/main.php?date=20221117

vate networks and public WAN to a central data processor node (COBALT) in the Netherlands. The LOFAR radio telescope enables astronomers to research, for example, the creation of the first stars and galaxies and cosmic magnetic structures [2].

Data Flow

A highly optimised data processing pipeline is required to efficiently convert analog cosmic signals to useful information. First, we capture the electromagnetic signals emitted by sources in the universe by antennas, e.g. with dish antennas or dipole antennas. The analogue signal is converted into the digital domain by FPGAs close to the antennas. These FPGAs perform the first of many filtering steps to reduce the data rate. The FPGAs also enable the division of the signal into different frequency bands for further processing. After this, the data can be sent to a central processing facility.

Central processing can be divided into three distinct stages: real-time processing, postprocessing, and archiving. In the first stage, real-time processing is required to create higher-quality data. During this stage, the samples from every station in the same frequency band are correlated, beamformed or a combination of both. This aims to detect statistical coherence in the received signals - otherwise too weak to be distinguished from the noise - and integrate the signals over a short period of time to reduce the output size. Beamforming creates one (or more) higher quality (thus better signal-to-noise ratio) signal(s). This is achieved by compensating the antenna signals for the phase shift between antennas whereafter they can be combined. This allows astronomers to look further into the universe at the cost of spatial resolution. The correlation method (also known as interferometry) applies cross-correlation to the antenna data. The result is a higher spatial resolution at the expense of sensitivity compared to beamforming. These operations are highly parallelisable and are therefore performed on GPUs.

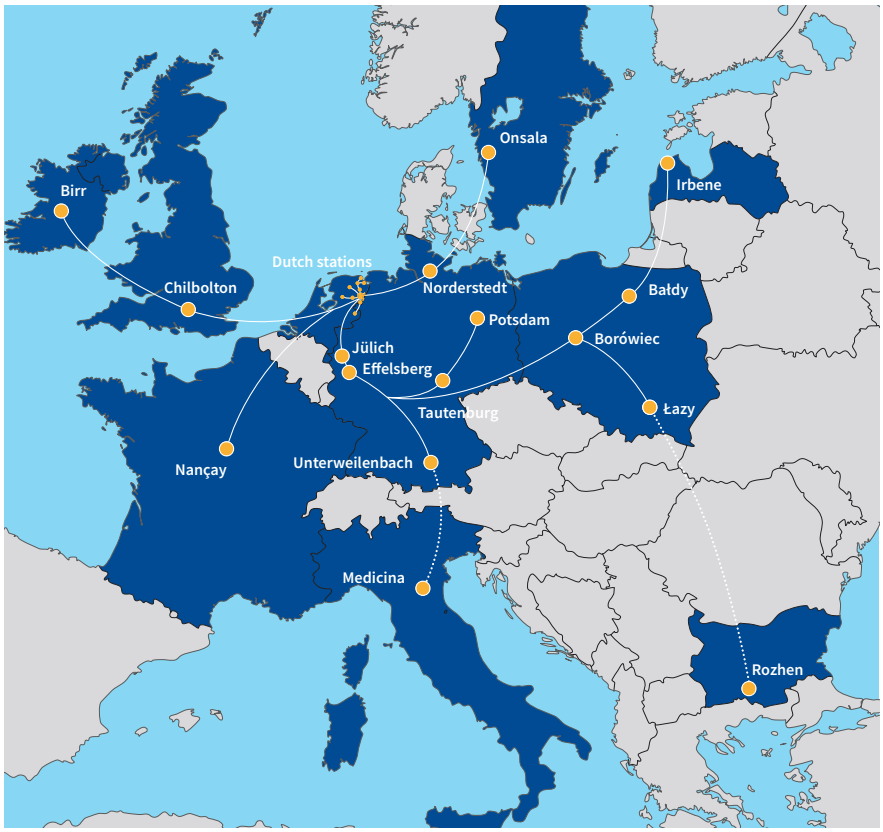

Figure 2. Map displaying the locations of LOFAR stations in Europe [1].

In the second stage, postprocessing is performed to calibrate the instruments, identify and remove possible interference and, lastly, create the final products (images, pulse profiles, source lists). Since the data rate is already significantly reduced after real-time processing, data from intermediate products during post-processing can be stored in DRAM or disks for longer periods. In the third and last stage, the final data products are sent to an archive and made available to astronomers for their research.

The data challenge

State-of-the-art radio telescopes produce a lot of data since the radio signals originating from space are measured with a large number of sensors (20-10000+), with sensors operating at very high sample rates (20 MHz up to 70 GHz), or a combination of both. Due to physical constraints and data arrangement, the processing cannot occur at the telescopes and requires a central facility where a supercomputer or computational cluster performs the calculations.

Such a system is cost-effective and flexible but requires a powerful computational back-end. Two important data transport characteristics in this application are: 1. the single direction in which the data is transported from the antenna to the computation facility, and 2. the static network and routing topology. The computational workload can be distributed over GPUs at a frequency band granularity, but each node requires the same frequency band from each antenna. Consequently, the data has to be transposed in the network. In addition, this data transport involves a long-term stable volume of data since an astronomical survey can take several hours or days.

Û

In this article, we focus on the data transport between the FPGA and the real-time processing on the GPU. Current transport implementations are based on UDP / IP transmission over (partially public) ethernet. On CPU nodes, this is implemented in a software stack that covers; a driver, a Linux kernel and user space. The data is copied multiple times between these layers to achieve separation and protection, consequently imposing a significant load on the CPU. Figure 4 shows the conventional way in which the FPGA communicates data to the GPU memory in the remote host. After the data is received in the user space, it can be moved to the GPU, again requiring interaction with the kernel stack. For large Ethernet data rates, this leads to significant inefficiency (or a bottleneck) in CPU utilisation. Current and near-future network technology will allow Ethernet line rates of up to 400Gbps, which are greatly beneficial for radio telescope systems. However, the current way of receiving and processing this (UDP-based) data is limited to 40 - 100Gbps.

Remote Direct Memory Access

In recent decades, Remote Direct Memory Access (RDMA) technologies have been developed to increase data throughput and improve latency while reducing CPU utilisation by offloading tasks to dedicated hardware, e.g. network cards. RDMA allows one host to access the memory of another host or device not in another physical system or chassis without CPU involvement in the data path (CPU bypassing). In addition to main (DRAM) memory, 3rd party memory (such as GPU memory) can also be used, such that data does not have to be stored temporarily in main memory. In other words, the data can be transferred directly into GPU memory, as illustrated in Figure 5. This way, the data traverses just once over the PCIe bus and does not pass the (software) kernel space several times. RDMA over Converged Ethernet version 2 (RoCEv2) is an RDMA protocol piggybacking onto standard Ethernet frames, IP and UDP for network routing. This implies that just the end nodes need to support the protocol while using a standard Ethernet network as an interconnect. RoCE is already supported in many network cards and widely used to transfer data between CPUs and GPUs, e.g. through NVIDIA GPUDirect [2].

However, RoCE has not been widely applied yet for systems with real-time streaming data with many senders, such as a radio telescope system. A proper FPGA implementation of RoCE and documentation on the impacts of the numerous RoCE settings for streaming data applications are missing. For example, the adoption of transport services (reliable or unreliable), sizes of the queues per connection, resource sharing between connections, the impact of memory accesses patterns and the number of memory channels.

Promising Results

A recent master's thesis study investigated whether RDMA over Converged Ethernet (RoCE) could be used for applications in radio telescopes [4]. First, a RoCE implementation for FPGAs was examined and tested. Key findings from this are that it is certainly feasible to use RoCE in an FPGA to write the data into DRAM main memory or GPU memory at high speed (90+ Gbps) and very low CPU utilisation (10% or less) in the receiving compute node.

In addition, through a representative set-up on the DAS6 [5] compute cluster, it was examined whether RoCE could manage the requirements for radio astronomical systems. Firstly, a naive comparison between UDP and RoCE demonstrated a 2x (50 to 97Gbps) and 3x (146% to 43%) improvement in the throughput and CPU utilisation when using RoCE, respectively. Yielding a significant 6x performance increase in just receiving the data.

As mentioned earlier, the RoCE protocol has many settings which can be optimised for different purposes, for example, throughput, latency and CPU utilisation. A setup taking the scale of the LOFAR system as a reference point concluded that RoCE can provide more efficient data transport. In this case, the benchmarking application used 2000 connections (in 4-to-1 topology), which could receive 90Gbps in GPU memory with 40% CPU utilisation of a single core. This data rate can be fur-

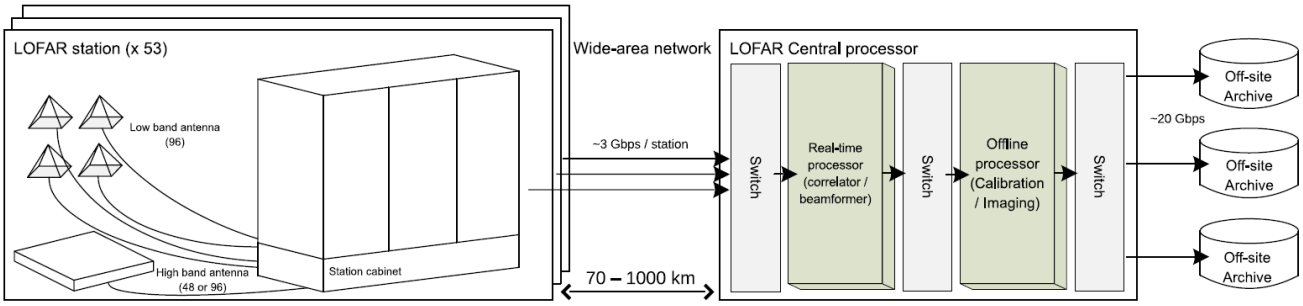

Figure 3. Top level view of the LOFAR front-end and back-end from antenna to off-site archive (Credit: P. Chris Broekema et. al. [3]).

[2] https://developer.nvidia.com/gpudirect

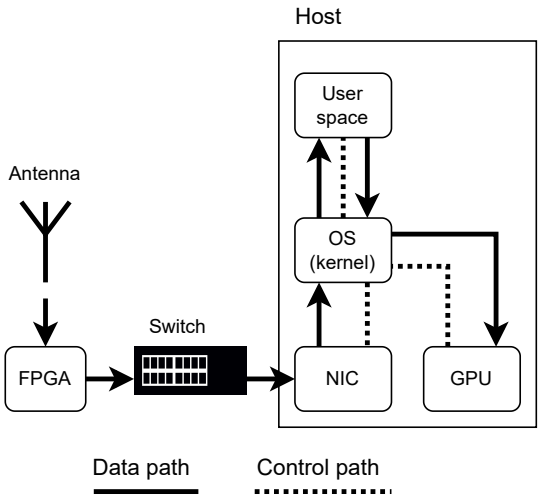

Figure 4. Vanilla data transport through software stack.

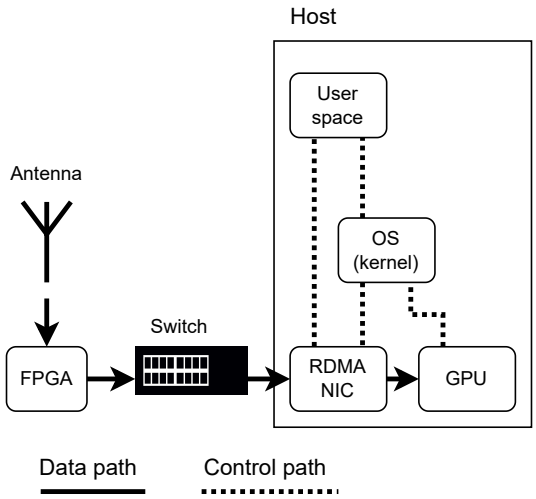

Figure 5. RDMA data transport with zero-copy.

ther increased by better congestion management, and CPU utilisation can be (drastically) reduced by sending more data per operation.

However, there is a trade-off in the amount of data per operation. The loss of one Ethernet frame can lead to the loss of all data belonging to the operation due to the selection of an unreliable connection. On the other hand, a small operation size (4kB, 8kB) with many connections (2000+) can lead to throughput degradation due to the high-speed scattered memory access patterns. Because insight into lost data is needed, the received data must be tracked per connection. As the number of connections increases, the alternation between connection information increases. Because of the high speed and low latency characteristics, ensuring this information stays close to the CPU (in cache) and can be accessed efficiently is advantageous.

Outlook

With the increasing deployment of hybrid computing (using CPUs, GPUs, FPGAs, and other dedicated hardware for dedicated compute applications), it becomes increasingly important to transfer data efficiently between the compute resources. An FPGA implementation that is compliant with the RoCEv2 standard will be beneficial for many applications that face large volumes of streaming data. RoCE is able to provide high throughput, low latency and low CPU overheads. Yet, to fully exploit its strengths, it needs to be supported by other components in the system, and the settings must be fine-tuned for the respective application.

The master's thesis study has shown that RoCEv2 is a promising technology: 1) to enable data streaming over commodity Ethernet beyond 100Gbps bandwidths in a radio telescope system; 2) to significantly increase the energy efficiency of current systems by decreasing CPU load; 3) as well as a combination of both.

Ê

[1] ASTRON. LOFAR telescope website. url: https://www.astron.nl/telescopes/lofar.
[2] Jean-Mathias Grießmeier, Philippe Zarka, and Michel Tagger. "Radioastronomy with LOFAR". In: ComptesRendus Physique 13.1 (2012), pp. 23–27. issn: 1631-0705. doi: https://doi.org/10.1016/j.crhy.2011.11.002. url: https://www.sciencedirect.com/science/article/pii/S1631070511002234.
[3] P. Chris Broekema et al. "Cobalt: A GPU-based correlator and beamformer for LOFAR". In: Astronomy and Computing 23 (2018), pp. 180–192. issn: 2213-1337. doi: https://doi.org/10.1016/j.ascom.2018.04.006. url: https://www.sciencedirect.com/science/article/pii/S2213133717301439.
[4] Willem de Laat. The Application of RDMA over Converged Ethernet Data Transport for Radio-Astronomy Systems. Dec. 2022. url: http://resolver.tudelft.nl/uuid:94d16d94-e4c7-4a54-a26a-af114db02dd7.
[5] Henri Bal et al. "A Medium-Scale Distributed System for Computer Science Research: Infrastructure for the Long Term". In: Computer 49.5 (2016), pp. 54–63. doi: 10.1109/MC.2016.127.

Energy Harvesting from Vibration

The First Inductor-less Bias-Flip Rectifier for Piezoelectric Energy Harvesting

By Sijun Du

Piezoelectric vibration energy harvesters have drawn much interest for powering self-sustained electronic devices. Furthermore, the continuous push toward miniaturization and higher levels of integration continues to form key drivers for autonomous sensor systems being developed as parts of the emerging Internet of Things (IoT) paradigm. An inductorless bias-flip rectifier is proposed in this article to perform residual charge inversion using capacitors. The voltage flip efficiency goes up to 80% while eight switched capacitors are employed[1]. The proposed SSH on capacitors circuit is designed and fabricated in a 0.35-µm CMOS process. This performance improvement is higher than most of the reported state-of-the-art inductor-based interface circuits, while the proposed circuit has a significantly smaller overall volume enabling system miniaturization.

Introduction

Internet-of-Things (IoT) is making our lives and work easier by connecting the physical world to the Internet. Due to the significant number of IoT nodes implemented ubiquitously in the world, using batteries to power them is obviously not an environmentally friendly or economical way. Driven by this, researchers started to think of a solution to let the IoT nodes be self-sustained with environmental energy. To make this happen, energy harvesting technology became a hot topic in both academia and industry.

Piezoelectric vibration energy harvesting (PVEH) has drawn much interest in recent years as a means of harvesting ambient kinetic energy to power wireless sensors and portable and implantable electronics. Among the various candidate vibration energy harvesting techniques considered, piezoelectric materials are widely used due to their relatively high power density, scalability and compatibility with conventional integrated circuit technologies. As the energy generated by a piezoelectric transducer (PT) cannot be directly used to power load electronics, an interface circuit is needed to rectify the output power and provide a stable supply. A typical piezoelectric VEH can provide a power density of around 10 - 500 µW/cm^{3}g^2, which sets a significant constraint on designing the associated power conditioning interface circuit. Full-bridge rectifiers (FBR) are widely used due to their simplicity and stability; however, their power efficiencies are relatively low as they set high voltage thresholds for the input voltage to overcome prior to any energy extraction.

Around 2016, various interface circuits have been reported based on inductors employed to improve the power efficiency with RLC loops. The SSHI (Synchronized Switch Harvesting on Inductor) rectifier (or inductor-based bias-flip) is one of the most energy-efficient circuits with ideal no-charge wastage developed for this purpose, which synchronously flips the voltage across the PT to minimize energy wastage due to charging the internal capacitor. However, most of these reported circuits require large inductors, up to 10 mH, to achieve acceptable efficiencies and these large inductors significantly increase the system volume, counter to the requirement for miniaturized self-powered systems.

In this article, an alternative SSHC (Synchronized Switch Harvesting on Capacitors) approach is presented to synchronously flip the voltage across

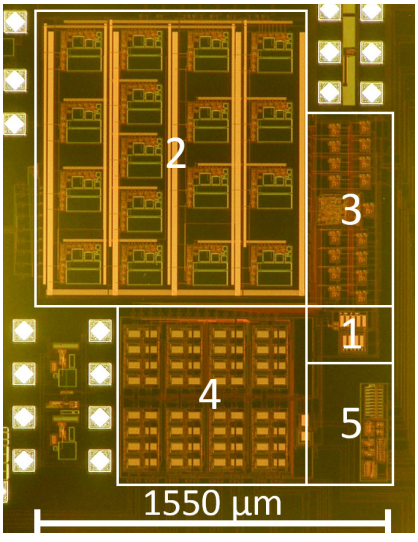

Figure 1. Sijun Du and his miniatrurized SSHC interface.

the PT using one or multiple switched capacitors instead of an inductor. This design does not require any inductor, thus significantly reducing the required system volume. This feature is particularly necessary when considering the design of miniaturized energy harvesting systems integrating MEMS (microelectromechanical systems) harvesters for applications to implantable devices and miniaturized wireless sensor nodes. Compared to reported state-of-the-art interface circuits, the proposed circuit also achieves higher voltage flip efficiency, hence higher energy extraction efficiency

Inductor-based SSHI nterface

Figure 2a shows the circuit schematic of a parallel-SSHI rectifier, which consists of a full-bridge rectifier (FBR) with a switch-controlled inductor to synchronously flip the voltage across the PT. A weakly coupled piezoelectric transducer (PT) is employed in this work; hence, the synchronized switch damping (SSD) effect [18] is neglected and the PT can be modeled as a current source I_P in parallel with a capacitor C_P [9]. The associated waveforms of the SSHI circuit are shown in Figure 2b. Before zero-crossing instants of the current source I_P, the voltage across the PT, V_{PT}, equals to V_S+2V_D or $-(V_S+2V_D)$. In order to overcome the threshold set by the FBR and transfer energy into the storage capacitor C_S in the next half-cycle, V_{PT} needs to be flipped from V_S+2V_D to $-(V_S+2V_D)$ (or vice-versa).

In an SSHI interface, analog switches driven by a synchronized pulse signal φ_{SSHI} are employed to control the RLC oscillation loop to flip the voltage. The resulting flipped voltage V_F is always lower than V_S+2V_D due to the resistive damping in the RLC loop, which can be written as:

$$V_F = (V_S + 2V_D)e^{-\frac{\pi}{\sqrt{\frac{4L}{R^2C}-1}}}$$

After the voltage flip, $|V_{PT}|$ needs to be charged from V_F to V_S+2V_D and this amount of energy is wasted. Therefore, the power efficiency of an SSHI interface usually depends on the voltage flip efficiency, which is expressed as:

$$\eta_{SSHI} = \frac{V_F}{V_S+2V_D} = e^{-\frac{\pi}{\sqrt{\frac{4L}{R^2C_P}-1}}}$$

where C_P, L and R represent the internal capacitor of the PT, the inductor and total resistance in the RLC loop, respectively. As C_P is inherently constant for a given PT, η_{SSHI} can only be increased by increasing L or decreasing R. In order to miniaturize the system, L is typically chosen in the range of a few mH; however, an inductor of this value still occupies significant system volume. While decreasing R, the contributory factors of R should be mentioned.

Usually, R is the total resistance in the RLC circuit, which includes DC resistance of the inductor R_{IND}, ON resistance of CMOS switches R_{SW} and other parasitic resistance R_{PAR}, such as CMOS wiring, vias and contacts. In terms of R_{IND}, it is usually proportional to the inductance L for a given inductor family. Hence, R_{IND} should also be considered while choosing the inductor for SSHI circuits. In order to reduce R_{SW}, the transistor sizes of the two analog switches shown in Figure 2a need to be designed to be larger; as a result, the gate capacitance of the transistors is increased. These large switches are usually power-hungry when driven. The following sections of this article propose a novel interface circuit with the ability of performing highly efficient voltage flipping without employing inductors, hence the energy efficiency is increased with smaller required volume.

Û

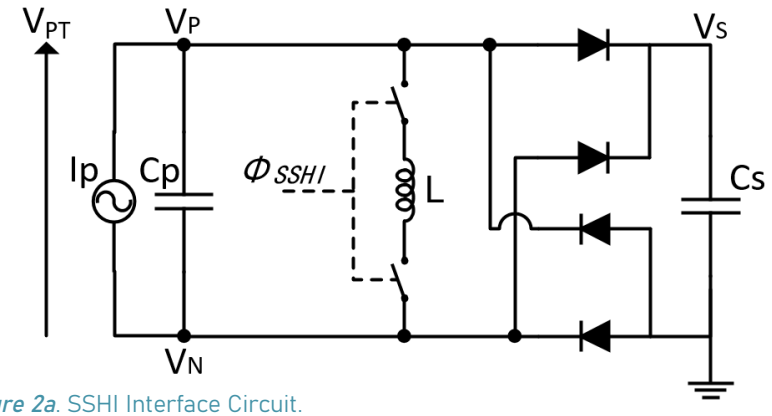

Figure 2a. SSHI Interface Circuit.

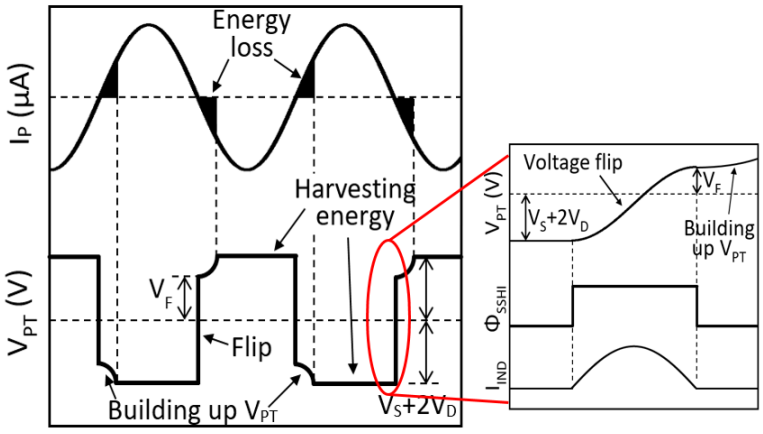

Figure 2b. Associated Waveform of SSHI Interface.

An SSHC Interface Circuit

In this section, an inductor-less inteface circuit is introduced, which employs one or multiple synchronized switched capacitors (SCs) to increase voltage flip efficiency and hence power extraction efficiency. The performance is then compared with an SSHI interface.

SSHC with one capacitor

Figure 3a shows the circuit diagram of the proposed SSHC (Synchronized Switch Harvesting on Capacitors) interface circuit with one switched capacitor C_1, or it can be called a charge-swap capacitor. In order to perform the charge inversion, five analogue switches driven by three pulse signals (φ_p, φ_0 and φ_n)are used. The three non-overlapping switching signals are synchronously generated to turn on the five switches sequentially in a specific order. The order of the three pulses depends on the polarization of the voltage V_{PT}.

Figure 3b shows the waveforms of the voltage V_{PT}, the voltage across the capacitor C_1 and the three pulse signals driving the five switches. At each zero-crossing moment of I_P, the three pulse signals (φ_p, φ_0 and φ_n) are sequentially generated to flip the voltage V_{PT}. Assuming $V_{PT} = V_S + 2V_D$ before the flipping instant (the left zoom-in figure), V_{PT} needs to be flipped towards negative. In this case, the pulse φ_p is first generated to damp a part of charge from C_P into the charge-swap capacitor C_1.

Then, the pulse φ_0 clears the residual charge in C_P and the pulse φn charges C_P from C_1 in the opposite sense. This allows the voltage V_{PT} to be partially flipped. While V_{PT} is supposed to be flipped from $-(V_S + 2V_D)$ towards positive polarity, the three pulses are now generated in an inversed order: φn to φ0 to φp (the right zoom-in figure).

As shown in the figure, the voltage flip efficiency is around 1/3. This is the optimal flip efficiency while using one charge-swap capacitor and the theoretical discussion can be found in [1].

SSHC with multiple capacitors

In order to flip additional charge across the capacitor C_P , more synchronized switched capacitors can be added to transfer more charge from C_P into a series of capacitors and conversely charge C_P to a higher voltage level. Figure 4a shows the proposed SSHC interface circuit with k switched capacitors. In this design, there are 4k+1 analog switches and 2k+1 switching signal phases: φ_0, φ_{1p}, φ_{1n}, φ_{2p}, φ_{2n}, φ_{3p}, φ_{3n}, etc. Assuming the number of switched capacitors is k = 8, the instant when VPT is being flipped from $-(V_S+2V_D)$ towards positive and from $V_S + 2V_D$ towards negative are shown in Figure 4b. The voltage V_{PT} and the 17 (as 2k+1 while k=8) phases of the switching signals are shown in the figure. From the figure, it can be seen that, in order to flip V_{PT} from $-(V_S+2V_D)$ towards positive, the phase order of the 17 pulses is: φ_{1n}, ..., φ_{8n}, φ_0, φ_{8p}, ..., φ_{1p}.

The first 8 phases aim to sequentially transfer charge from C_P to the 8 switched capacitors, C_1 to C_8. The phase φ_0 clears the residual charge in C_P and the following 8 phases sequentially connect the 8 switched capacitors in an opposite sense to flip the voltage V_{PT} . While V_{PT} needs to be flipped from V_S+2V_D towards negative, the phase order of the 17 pulses is completely reversed, as shown in the figure. The theoretical discussion on optimal voltage flip efficiencies is omitted here, which can be found in [1]. As a partial conclusion, the theoretical voltage filp efficiency can achieve 80%.

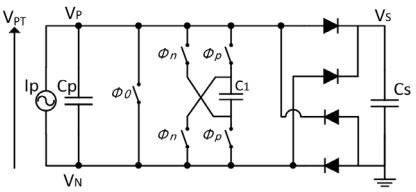

Figure 3a. SSHC interface with one charge-swap capacitor.

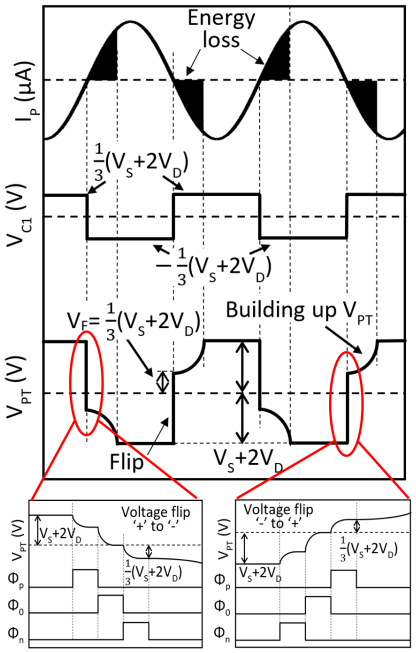

Figure 3b. Proposed SSHC interface circuit and the associated waveforms.

Measurement Results

The proposed SSHC interface circuit was designed and fabricated in a 0.35 µm HV CMOS process. The system was experimentally evaluated using a commercially available piezoelectric transducer (PT) of dimension 58mm × 16mm (Mide Technology Corporation V21BL). Figure 5 shows the output power of the PT with a convention-

About the Author

Sijun Du joined the Electronic Instrumentation Laboratory, Department of Microelectronics, TU Delft in 2020, where he is now an Assistant Professor (tenured).

His current research group is focused on energy-efficient integrated circuits and systems, including energy harvesting, wireless power transfer, and DC/DC converters used in autonomous wireless sensors for the Internet of Things (IoT), wearable electronics, biomedical devices, and microrobots.

al full bridge rectifier (FBR) and with the proposed SSHC rectifier with one switched capacitor and eight switched capacitors. They are measured with V_S fixed to 5V and the open circuit voltage level V_{OS}, which is the RMS value of supplied VPT, varying from 0 to 15 V. In the figure, it can be seen that the proposed SSHC interface with eight SCs can be functional under a smaller VOC and larger output power than either SSHC with two SCs or FBR. The proposed SSHC interface with eight SCs can provide output power up to 1.2 mW.

Conclusion

This article introduced an inductor-less interface circuit for piezoelectric vibration-based energy harvesters employing switched capacitors to synchronously flip the residual charge across the piezoelectric transducer (PT) to significantly improve key circuit metrics. Compared to reported state-of-the-art interface circuits, such as SSHI (synchronized switch harvesting on inductor) and other approaches, the proposed interface circuit completely removes the requirement for an inductor to flip the voltage across the PT. With theoretical calculations, the voltage flip efficiency achieves 80% with 8 switched capacitors. For an SSHI interface circuit to achieve equal voltage flip efficiency, a large inductor is required, which is very impractical in miniaturized systems for real-world implementations. The measured results show that the proposed SSHC interface circuit can output a stable output power up to 1.2 mW, which is sufficient to power most of wireless IoT devices. A follow-up work with all the switched capacitors integrated on-chip enables a new-class of fully integrated self-powered CMOS-MEMS sensor nodes, which was presented in [2].

Ê

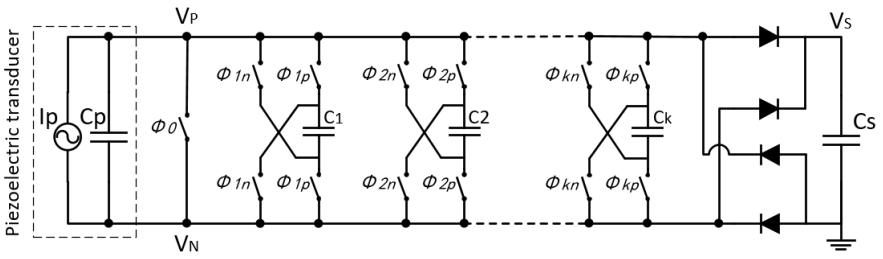

Figure 4a. Proposed SSHC interface circuit with k synchronized switched capacitors.

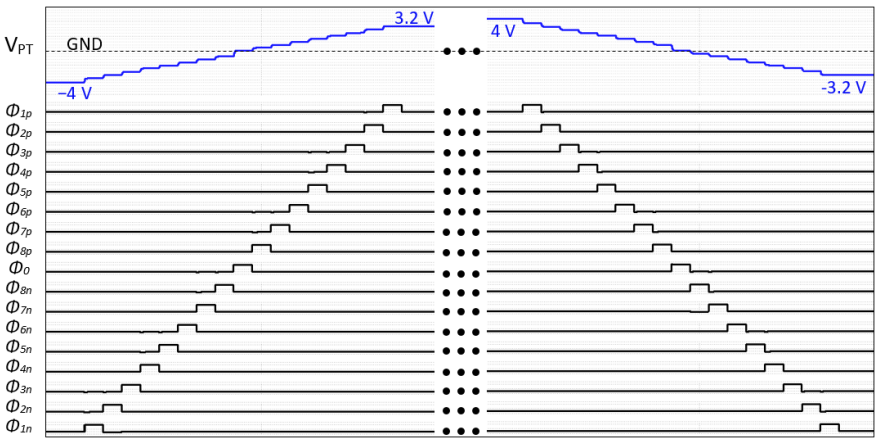

Figure 4b. Associated waveform of SSHC interface with k synchronized switched capacitors.

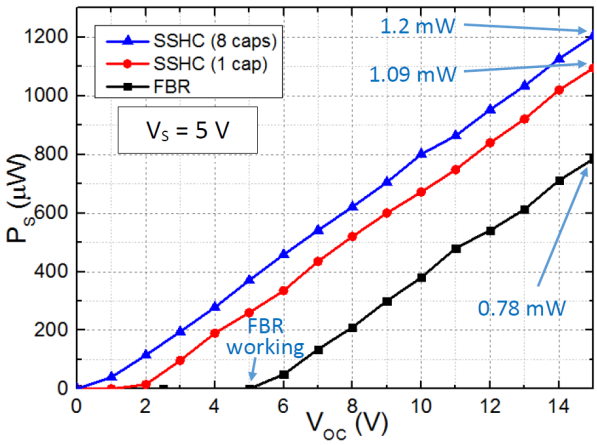

Figure 5. Measurement results of P_S-V_{OC} graph with Vs=5V.

[1] S. Du and A. A. Seshia, "An Inductorless Bias-Flip Rectifier for Piezoelectric Energy Harvesting," in IEEE Journal of Solid-State Circuits, vol. 52, no. 10, pp. 2746-2757, Oct. 2017, doi: 10.1109/JSSC.2017.2725959.

[2] S. Du and A. A. Seshia, "A fully integrated split-electrode synchronized-switch-harvesting-on-capacitors (SE-SSHC) rectifier for piezoelectric energy harvesting with between 358% and 821% power-extraction enhancement," in 2018 IEEE International Solid - State Circuits Conference - (ISSCC), 11-15 Feb. 2018

Advertorial

Polymetallic Nodule Collection at Allseas

By Frank Teunisse, Unit Head E&I at Allseas in Delft

Polymetallic nodules, also called manganese nodules, were discovered by ocean explorers in the 19th century. Even then, scientists on board were speculating about the potential mineral wealth inside these potato-sized rock formations. Lying on the seafloor unattached, these nodules contain high grades of cobalt, copper, nickel and manganese – metals required to manufacture batteries for storing renewable energy and powering electric vehicles.
For decades, these mineral-rich resources remained unexploited due to the lack of viable technology and the absence of a governing body to oversee activities in the deep ocean and international waters. More recently, advances in extraction technology, falling yields from land-based sources and growing demand for greener technologies have all served to push the deep-sea mineral collection higher up the agenda.

Seafloor to shore

Allseas has leveraged its deep-water expertise and developed technology to responsibly recover polymetallic nodules from the ocean floor and transfer them to the surface for transport to shore for processing. We have engineered an entire seafloor-to-shore supply chain. At the heart of the operation is the subsea collector vehicle, operated from our surface production vessel, Hidden Gem, and connected by a vertical transport system (VTS).

Built at 20% commercial scale, the prototype collector measures 12 metres long, 6 metres wide and 5 metres high, and has the capacity to collect over 80 tonnes of nodules per hour. For optimal manoeuvrability in the soft soil conditions, we focused on weight during the detail design phase. A lightweight frame, high-density plastic for processing equipment and tailored lightweight track design inspired by ski resort snow groomers. The key feature, hydraulic pick-up nozzles, is designed to optimise pick-up efficiency and minimise sediment disturbance, so-called "plumes".

Polymetallic nodules tyically vary between 5 and 10 cm in diameter, similar to the size of potatoes.

In parallel, we designed an air-lift system for pumping nodules to the surface via a 4.2-kilometre production riser. To generate the required uplift, air is injected at intense pressure about halfway up the riser. The collector is attached to the riser via a 500-metre flexible hose, which provides the necessary freedom during operation.

Electrical challenges

It is vital that the collector continues to operate 24-7 while withstanding the extreme pressures at ultra-water depths. This presents its own electrical challenges when designing the power and control systems. The collector requires more than 1.2 megawatts of power, and that power must travel 5 kilometres from surface to seabed.

The Hidden Gem.

Intense moments in the control room.

We found a partner for custom-designing a 6.6 kV "umbilical" cable with 6 independent 3 phase feeders (4x power, 2x control) and fibres for data connections. We used frequency converters to reduce the inrush currents of the motors, allowing smaller conductor sizes and thus a smaller and lighter umbilical. The umbilical must be strong enough to support its own weight during deployment and lift the collector during recovery. For ultimate strength, more than half is made from aramid fibres (same material as Kevlar®).

On the collector itself, the power conductors are directly wired to the electric motors. The control feeders and fibres are connected to several "control pods", metal canisters housing cameras, sensors, valves, and actuators. As the required wall thickness of the canisters depends on the internal volume, the electronics had to be designed very space efficient.

Putting it to the test

After an intense 3-year development cycle, in September 2022 it was time to put everything to the test. A 130-strong team comprising offshore crew and engineers from the office departed on a 2-month test mission in the remote Pacific Ocean between Hawaii and Mexico. It was an enormous success.

Our state-of-the-art system achieved all ambitious commissioning and production milestones – and sustained production rate of 86.4 tonnes per hour – during the first fully integrated deep-ocean nodule collection tests since the 1970s. In doing so, Allseas broke new ground, rounding off the trials with a record haul of more than 3000 tonnes of nodules. The sight of the first nodules arriving on board was met with huge excitement!

The trials also presented a unique opportunity to dozens of marine experts and scientists to conduct extensive environmental monitoring and impact research. Insight gained from the data will offer opportunity to advance the design further.

We are now making plans for a commercial-size collector and VTS, targeting a production capacity of 1.3 million tonnes of wet nodules per year. We expect production readiness by the end of 2024. Parallel, we are also looking at systems for handling, offloading and transporting the nodules to land for processing. Execution of these plans rests on the outcome of decision by the International Seabed Authority (ISA), the regulatory body, on regulations governing deep-sea mining and the award of commercial exploitation licenses.

These are busy times, but it's great to be part of a team laying the foundations for another new, exciting chapter in Allseas' history of pushing technical boundaries.

Joining Allseas

Are you also not afraid of adventure and to go where no-one went before? Come pioneer with us and contact me at electrical@allseas.com or find our vacancies on www.Allseas.com/careers.

Ê

Linkedin.com/company/Allseas/
@AllseasPioneers
Facebook.com/Allseaspioneers

First nodules arriving on deck.

Crew inspect the collector ahead of deployment.

Study Collection

The unveiled history of governmental message security (1945 - 1960)

By M.R. Oberman MSc

This article is motivated by the appearance of a book describing post-war developments in cryptography[1] and the recent acquisition of a rare Ecolex IV cryptographic machine by the Historic Collection of EEMCS in Delft. This "story" is about the security of electronic messaging communication. It started right after World War II when the telex network was the only carrier for electronic message traffic from desk to desk. The telex of that time was the precursor of our electronic mail and the follow-up of telegrams transmitted by Morse codes at the post office in the 1930s.

Telex is an abbreviation for teleprinter exchange. In a telex network, two tele printers send each other messages through a circuit-switched network. A teleprinter (Figure 1) can be directly operated by a person or can read character-based information from a paper tape reader. A printing mechanism prints the text on the receiving end, and a tape punch may be used to produce a paper tape. In order to reduce connection time, messages were often pre-punched so that they could be transferred at the highest speed possible. The transmission protocol is the Baudot code where each character is encoded in 5 bits. These 5 bits are used to read/punch so-called 5-hole paper tape [1]. The code transmitted, sent via fixed or wireless networks at that time, was easily intercepted by eavesdroppers. This created a need for encryption.

As with many projects, there is a multitude of people involved. Of these, a few people stand out. One of those people was prof. dr. ir. R.M.M. Oberman, the father of the author of this paper. He worked as a research engineer at the Dutch PTT from 1936-1957. Later he became affiliated with Delft University of Technology (1947-1980). He was appointed as an Honorary Member of the ETV later on. Prof. Oberman was the driving force behind the development of the first post-war cryptographic equipment for the Dutch government. The funding for research on electronic messages security after WWII was supplied by the Ministry of Foreign Affairs in The Hague, needing secret communications between the ministry and some of the most important Dutch embassies. Similar crypto-equipment was also used at that time (1946-1957) on all Dutch Royal Navy ships.

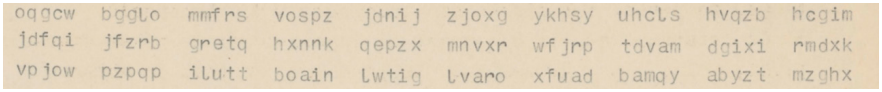

Figure 2. The formatting of a cipher text.

Oberman was well known for his developments in public switching technology. The same technique led to cryptography implementations, which immediately became a classified issue because of its governmental applications. This is the reason why nothing was ever written about it until I got my hands on his personal- and classified archive!

During WWII, every nation put emphasis on trying to read and decode messages from armies and diplomatic services from other countries, even from their allies. It became apparent that cryptographic equipment as used by several embassies was rather vulnerable to code-cracking. Crypto equipment was seen after WWII as the only mechanism to ensure fully secure electronic messages. Besides the added security it can be seen as the first step in administrative automation. An electronic cryptosystem replaced several people in a code room where coding and decoding messages was the main task. Through this development, fewer people had knowledge of classified information.

Figure 1. The Siemens Telex machine.

OTP

All crypto equipment immediately after WWII was based on the one-time pad system (OTP) [2, 3] which was the only system of that time that was proven to be unbreakable. Even in the age of quantum computing, this is to be expected. OTP stands for One Time Pad and is based on the use of a one-time, randomly composed key. In essence, a structured message (i.e. plain text) is drowned in real noise. The mixed sum of plain text and the random key is called the cypher text. Since the plain text is drowned in the noise, it is invisible and therefore unreadable. When

[1] www.oberman.nl/boek

Figure 3. The technician M. Koppenberg in front of "his" Colex.

the same noise is removed exactly from the received text, position by position with the same value (key-value) as on the sending side, the plain text reappears. The key generator was basically a perfect white noise generator. White noise as a technical item was well known to PTT being related to the simulation of telephone traffic to test switches.

One problem with OTP is the distribution of the key. The key must be available on both the sending and receiving sides before communications can be started. How do you ensure that the key text is available on both sides? A white noise OTP key generator cannot generate the same key text again. To overcome this problem, the white noise was recorded as a sequence of truly random numbers and punched on a paper tape (the key tape). The exact copy was made by a tape duplicator. The Ministry of Foreign Affairs in The Hague used its well-organized diplomatic postal system for the key distribution. For the Royal Dutch Navy, tapes were supplied during the ship visits to the home port. Proper key management was the crux behind the strict OTP security requirements.
OTP systems are symmetrical cryptosystems. They have identical keys on both sides, sending and receiving. OTP was the standard until 1975 when asymmetrical systems were invented. The sender and receiver have a different but mathematically related connection in asymmetric systems. This also laid the foundation for Bitcoin and other cryptocurrencies.

Colex

The Colex (Code-telex) system, like any OTP system, consists of two parts: the key generator and the mixer.
The key text must be at least as long as the message it protects and must comply with the white noise properties, i.e. be completely structureless and not be reproducible with the key generator. At that time it was a technical miracle. An example: a rotating drum suddenly stopped will give data with white noise properties. In the Ecolex systems, an advanced control mechanism was used to check the white noise statistics to be sure the key is fully random.
The mixer combines the key text with the message (plain text) producing the cypher text to be sent via the network. Such a mixer is network-dependent. When the network uses in-band signalling, the cypher text must not contain any control characters. At that time, the cypher text message format had a universal character: It consisted of a number of lines; each line is composed of 10 groups of 5 characters followed by two spaces (Figure 2).

A working system

Oberman had studied electrical engineering, as well as mechanical engineering at the then TH-Delft. This was special and practical for the development of the encryption system in those days, because in addition to the knowledge of switching properties and the needed electronic components, he also had knowledge of and insight into the mechanical properties of the system. E.g., the system must be able to work without maintenance or be locally repairable. After all, the systems could be located everywhere in the world in environments where people had little knowledge of this technology. It was almost impossible to service a defective system at remote locations such as embassies worldwide.

It takes more than just telex knowledge to make a good encryption device. In addition, there are several very different elements that make a system good. For example it had to be well

Û

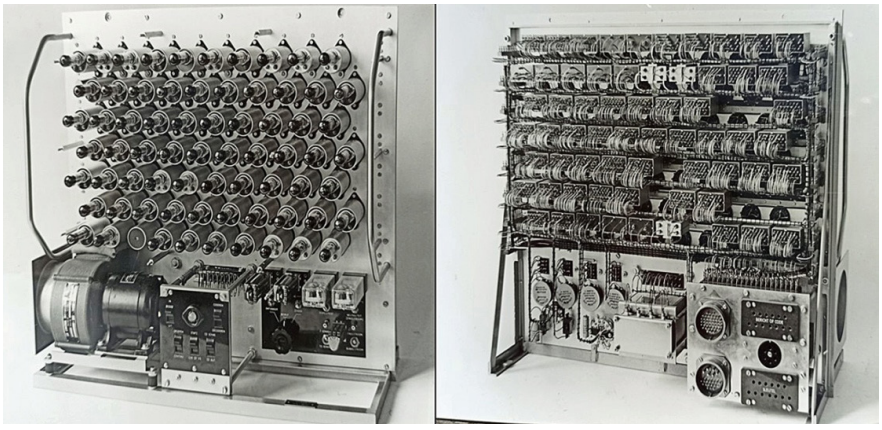

Figure 4. The Ecolex using radio tubes (front view and back view).

known within PTT in the same way as telex because of a diversity of subjects like functional standards to follow, tariffing, physical connection, network signalling and management issues. Encryption had to fit seamlessly into the worldwide telex network. Other important design questions were:

- What is a strong cryptosystem and why?
- How do we build a robust 24/7 running system?
- How to build a low-(non?) maintenance system? PTT had no service organization outside the national country's borders.
- How to build a system where mistakes do not lead to the release of the key text or the plain text?

Prof. Oberman first constructed a version of the Colex using relays (Figure 3). Other components were simply not available during the first post-war years. This relay-driven Colex encrypted 3 to 6 characters per second. In 1950 a radio tube follow-up was made – a functional 1-to-1 replacement - encrypting 40 characters per second, called the Ecolex (Figure 4).
The Colex was used from mid-1948 and officially put into service on April 5th 1949 by prime minister W. Drees in The Hague. At that time, it connected the embassies of Paris, London, Washington and Jakarta with the Foreign Affairs Office in The Hague. Six systems were made, each with about 98 identical relays. The switch to radio tubes made the relays intended for the Colex not usable anymore. However, the Colex relays had been purchased for twelve systems. The result: about six hundred relays had become redundant.

Transistors

Through visits to the USA in the course of 1952, including IBM and Westinghouse, Oberman had concluded that radio tubes were past their prime time. The transistor was, in his opinion, the better technical invention to replace radio tubes as the basic switching component. In his opinion, the transistor was the component that would take a fundamental step towards the realization of automation systems, computer systems, and also the construction of cryptosystems. Now we do know that the transistor became a future-proof solution, but in those days the considerations were:

- The basic switchgear unit was orders of magnitude smaller; the result: less extensive systems.
- Lower energy consumption, resulting in less heat generation.
- Faster switching capabilities.
- More robust than radio tubes, which means less maintenance.
- Non-hazardous low voltage and therefore more maintenance-safe.

Philips

In 1956 Oberman made the migration with his cryptosystems from the radio tube-based Ecolex to a system based on transistors as the core switching component. This was only 10 years after the first relay system was made by him and his team. Unfortunately, PTT wanted to eliminate everything related to the topic of cryptography. That change in policy was made a few years before 1957 when Oberman and his coworker Snijders [2] left PTT. Cryptography and everything related to it was passed on from PTT to Philips USFA in Eindhoven. In this way, USFA could make a quick start in the crypto field. Philips USFA produced the Ecolex IV of which a specimen was recently added to the Historic Collection in Delft (Figure 5). In 2003, the successor to Philips USFA stopped making crypto equipment in the Netherlands. A successful story came to an end.
The Colex and the Ecolex were simplex OTP systems. During WWII the theory behind the operation of OTP systems was already more than 35 years old and therefore widely known. The mathematical proof however was not given until WWII by Shannon [4]. There is no patent on OTP as a system since it is a well-known principle. A

What happened to the six hundred redundant Colex relays?
These relays were given by L. Kosten, who was head of the mathematical department at PTT in 1950, to the recently graduated Ir. W. L. van der Poel. Van der Poel joined PTT in 1950 immediately after graduating. His assignment was to build a computer. These redundant relays were the basis for the ARCO computer system which can still be seen in the Historic Collection.

Figure 5 . The Ecolex-IV, as recently acquired by the Historic Collection.

Figure 7. Prof. Oberman cycling at his department at the EWI-10th floor after his restart as extraordinary professor.

[2] A. Snijders cooperated closely with Oberman and also became a professor at TU Delft. He was the inventor of the state lottery system in 1970, which uses an OTP key generator.

patent could be obtained for a specific implementation thereof.
At that time, technology was very important to arrive at an effective solution. Additionally to technology, there must be a fit with the organizational environment. In short, this concerns the following points:

- The cryptographic work by prof. Oberman was quickly declared to be Top Secret classified information due to the sensitive type of application for the government. This sometimes led to less support or a lack of understanding of his work outside his immediate organizational environment.
- These struggles became the basis for his departure from PTT to the TU-Delft in January 1958.
- Prof. Oberman studied both mechanical and electrical engineering at the TU. This knowledge was crucial in order to make a good operational system in addition to a functional system.
- As a manager, little is known about Oberman historically. It was striking that when he left PTT for the TU, 10 out of the 13 people from his department at PTT followed him to the TU, to the switching department at the Electrical Engineering department at the TU-Delft within 2 years!

Besides these points the career of prof. Oberman has a few more remarkable aspects:
He holds more than 90 patents and he was the ultimate responsible person on behalf of the TU for the construction of the tall blue/orange EWI building, which was opened in 1969 and is still in use 52 years later.

Ê

Figure 8. M.R. Oberman, author.

[1] https://en.wikipedia.org/wiki/Baudot_code
[2] https://nl.wikipedia.org/wiki/One-time_pad
[3] www.cryptomuseum.com
[4] C. E. Shannon: Communication Theory of Secrecy Systems. Bell System Technical Journal, October 1949

IN MEMORIAM

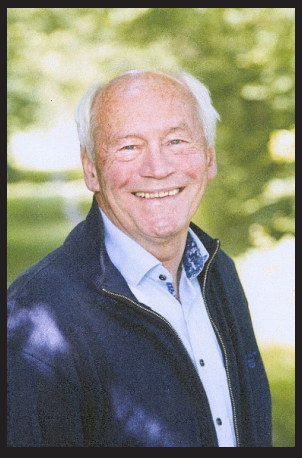

ir. R.A. (Rob) Timmermans

* March 29, 1947 September 2, 2022 †

Rob was active for 40 years in the areas of IT and telecommunication before devoting his time to the Study Collection. He specialised in the history of time distribution systems and the applications of time switches. Rob set up the exhibition “Op Tijd Schakelen”, which can be seen in the EEMCS low rise building.

Advertorial

True or false? Test your knowledge about ASML

From chipmaking to EUV and from the number of employees globally to next generation machines, discover the most important facts about our fascinating tech company.

The name 'ASML' is an acronym.

FALSE. ASML isn't an abbreviation of anything anymore, though it used to stand for 'Advanced Semiconductor Materials Lithography'. ASML was founded in 1984 as a joint venture between Philips and ASM International, so a name was chosen to reflect the partners in the venture. Over time, this name has become simply 'ASML'.

ASML makes microchips.

FALSE. ASML does not make microchips – we make the machines that other companies use to make microchips. We also don't make the silicon wafers that form the cradle of the chip. Customers such as Intel, Samsung and TSMC use ASML's DUV and EUV lithography systems to print tiny patterns on silicon that has been treated with 'photoresist' chemicals. They also rely on our metrology and inspection systems, together with our computational lithography and patterning control software solutions, to achieve the highest yield and best performance in mass production.

ASML is the only company that makes EUV (extreme ultraviolet) lithography technology.

TRUE. Unlike in the DUV (deep ultraviolet) lithography market, where ASML competes with other top-notch suppliers, ASML is currently the only lithography equipment supplier capable of producing EUV technology. Chipmakers use these EUV systems to manufacture the world's most advanced microchips. In fact, if you own a relatively new smartphone, gaming console or smart watch, chances are you've benefited directly from EUV lithography technology. We spent 20 years developing EUV with our partners and suppliers, resulting in a machine that contains around 100,000 parts. To ship just one of these huge machines to customers requires 40 freight containers, three cargo planes and 20 trucks.

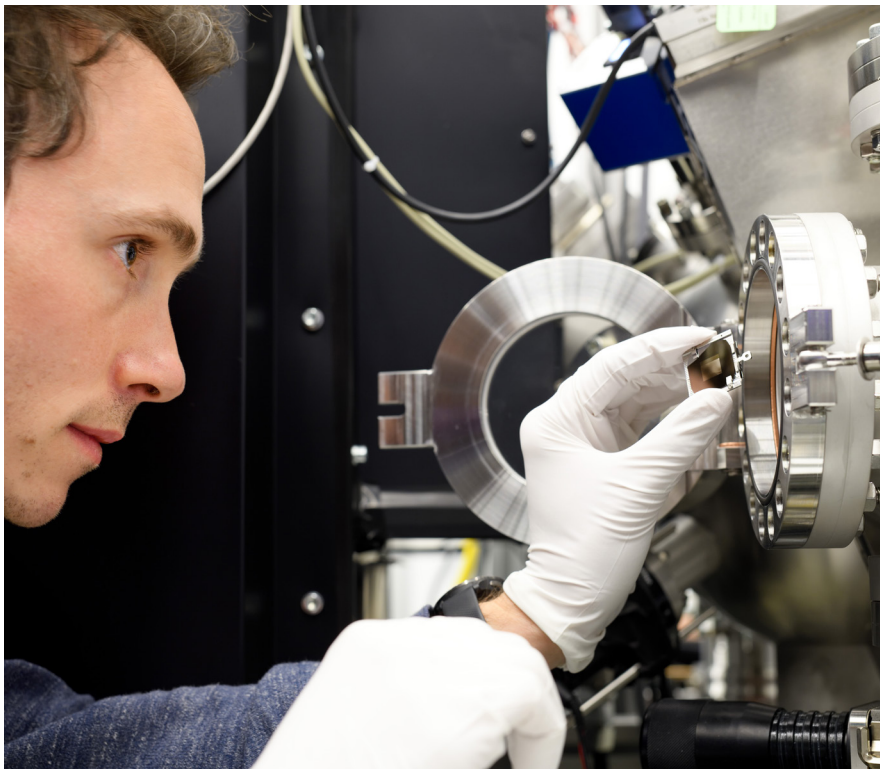

Figure 1. Sample preparation of a pellicle in the research lab.

Figure 2. Research on the EUV lithography systems in the cleanroom.

Figure 3. A team working on the assembly of an EUV system in the cleanroom.

An ASML machine is all you need to make microchips.

FALSE. Making chips is a complex, long and expensive process. Our customers have spent years and invested billions building 'fabs' (fabrication plants), buying equipment and training employees to become experts in the complex field of semiconductor manufacturing. ASML's lithography machines form an important part of a chipmaker's production line, but they are not all that's required to produce microchips. Lithography – printing patterns on silicon wafers – is certainly a critical step in the chipmaking process, but it's just one of many!

ASML is building a new kind of EUV lithography machine.

TRUE. In the semiconductor industry, innovation never stops. That's why we're already developing a next-generation EUV platform that increases the numerical aperture (NA) from 0.33 to 0.55. This means that the optics systems in the new machines will allow light with larger angles of incidence to hit the wafer, giving the system a higher resolution. The EUV 0.55 NA platform, called EXE, is well on its way to production – we're planning the first shipments of these machines to customers for R&D purposes by the end of 2023, and we expect them to be used in high-volume manufacturing by 2025.

At ASML, we're changemakers!

Our growing team of over 37,000 people and 144 nationalities provides leading chipmakers with the hardware, software and services to mass produce patterns on silicon. We're probably part of the device you use to communicate, learn or play games with.

Headquartered in Europe's prolific tech hub, the Brainport Eindhoven region in the Netherlands, we have over 60 locations globally and annual net sales of €18.6 billion in 2021.

Ê

Contact us

Curious to learn how you can be a part of progress? Contact our campus promoter Lieke van der Linden at your university at lieke@workingatasml.com with all of your questions about ASML or visit www.asml.com/students.

A Revised Bachelor EE Curriculum

A preview of the upcoming years by Electrical Engineerings director of studies

By N. van der Meijs

In the academic year 2023-'24, a revised Bachelor of Science in Electrical Engineering will start at TU Delft. This article aims to sketch the current situation.

Why is the curriculum revision necessary?

The current BSc EE curriculum's contents date back to 2010. One of its aims was to connect theory to practice, which led to developing the EPO's as a main means for that. EPO stands for Elektrotechniek Project Onderwijs (Electrical Engineering Project Education). Apart from a change in form in 2014 (three courses of five EC each quarter), the contents have not really been updated since and the programme is beginning to show its age. Much has changed in (electrical) engineering and its place in society. The words 'energy transition' did not dominate the news in 2010, and robotics, internet-of-things, bioelectronics, e-mobility and machine learning were mostly futuristic concepts. While the programme provides the basic theory for all of this, a revision is needed and modern topics must be added to make the curriculum (again) future-proof.

What are the principles behind the curriculum revision?

Something has to go if we want to make space for new topics. But since there is nothing really redundant, it is natural to introduce elective courses where students can select some courses based on their interests. These electives, and the curriculum as a whole, will connect to modern, relevant societal issues and better show the potential impact an electrical engineer can have in tackling the big problems of today and the future. Thus, the curriculum will show and better prepare students for topics such as mentioned above and more generally in fields like affordable and clean energy, health and well-being and sustainable cities and communities.

Apart from adding modern topics, other important goals for the revision include to improve the first-year retention ratio (increase the BSA percentage). For that, the programme should offer the right content on the right level in a motivating and activating way. It is known that motivation and engagement improve when students can have ownership of their own learning, where ownership in turn is fostered by allowing students to explore their interests and passions and choose directions that feel meaningful to them. This matches nicely with the idea of elective courses and topics.

Furthermore, it also turns out that students can benefit and feel more engaged and motivated when they have a feel of cohesion of the curriculum. Thus, we plan a course in the first quarter of the first year that aims to work as a gateway into Electrical Engineering. It will allow students to better understand how the courses fit into the overall programme, and how they come together to build all the knowledge and skills of a BSc engineer in Electrical Engineering. It will also show the students examples what EE is about and what problems and applications fields they will be able to work on after graduation; how the skills and knowledge they will gain will be relevant to those fields.

The revision also aims to attract more students, in response to the already years-long and future-projected shortage of electrical engineers entering the job market. EE-ers are highly needed to solve the scientific and industrial issues of today society, but there simply aren't enough of them. This might be true for more branches of engineering, but in EE the problem is particularly severe – the job-openings-to-graduates ratio is consistently topping the official statistics. We believe that the revised curriculum with clear connections to pressing societal issues will help prospective students to see the value of pursuing an EE degree and the opportunities they will have to make a difference in the world.

Last but not least, the curriculum revision aims to improve inclusivity and diversity, including gender balance. In a recent gender scan, the VHTO (Expert Centre Gender Diversity in STEM, engineering and IT) recognized the programme for its efforts to improve the gender balance. Yet clear opportunities for doing better still were sketched. It will help various groups of students to make the programme more transparent with frequent and explicit descriptions of the links between courses and the subject domain (as, e.g., done in the Gateway course). Such transparency will also help dedicated outreach activities. Without aiming to describe all aspects of inclusivity and diversity completely, it is good to note that the curriculum revision will also serve as a foundation for implementing optimized pedagogy.

This pedagogy aims to provide a more diverse range of study opportunities for students with varied learning preferences.

What does the revised curriculum look like?

A draft blueprint (module chart) of the curriculum is shown in Figure 1. This is close to but not final, the order of courses might still change a bit. Students will be informed about the final design sometime in February.

In this figure, the grey courses are compulsory for all students. Note that these courses are not one-to-one re-implementations of the current courses. To make space for the new course elements, the number of courses in some learning lines is reduced from the current programme. The Gateway course is visible in Q1 of the first year. Two elective courses are shown in Q3 of Y2 and of Y3. Among others, there will be elective courses on Electrical Energy and Networking. Note that both these courses are in a learning line that has a reduced number of courses but that these electives can more than compensate that for the students selecting those courses.
The 10 EC course in Q4 of Y2 will have elective variants. They will be called EE for Next Generation X, where X will for example be Energy Systems or Communication and Sensing. These courses will combine theory and project work and will offer rewarding views on how Electrical Engineering can help address future societal challenges.

IP 1 … IP 3 are coming in place of the EPO's; they are now named Integration Project to indicate their role in the curriculum more clearly – they integrate and apply the knowledge of the theory courses around them. There are only three IPs while there are four EPOs, but one such former EPO is absorbed by the Next Gen course.

How does the revision affect current students?

The new curriculum will start in the academic year 2023-2024 with a new first year, with the new second year starting one year later. The new third year can start again one year later or tentatively together with the new second year (this depends on choices still to be made).

Of course, program changes may affect the study path of students who do not follow the nominal schedule. However, several measures will be taken to mitigate the disadvantages. For example, all exams (first exams and resits) of a given year will be offered in the academic year after a course year is changed. Thus, students who are now in their first year and have missed one or more courses at the end of this year can still complete their original first-year program in the next academic year. This is partially facilitated by this year's lecture recordings. Nevertheless, we are also considering specific instructional aids (e.g., tutorials and Q&A sessions) when and where they can make a difference.

More generally, there will be rules and regulations governing how students can combine courses from the old and revised curriculum so that they can graduate without additional delay in a way that still ensures their level of education and training according to the program's final qualifications. Thus, if examinations of a course from the old curriculum are no longer offered, students will be able to replace that course with an appropriate (more or less equivalent) course from the new curriculum.

Overall, the transition rules can help to minimize the disruption and confusion caused by the curriculum revision. The academic counsellor and student mentors will help students navigate the changes and ensure that students are able to complete their studies and graduate as planned.

Summary

This article has briefly motivated the need for a curriculum revision and its main objectives and has shown an unfinished draft of the curriculum blueprint. Finally, it has explained how transition rules will mitigate potentially adverse effects on study planning for students who are not on a BSc-in-3 schedule. We look forward to an enticing new curriculum next year!

Ê

	1st Quarter	2nd Quarter	3rd Quarter	4th Quarter
Year 1	Math 1	Math 2	Math 3	Energy
	Circuits 1	Circuits 2	Digital 1	Digital 2
	Gateway into EE	IP 1	Physics 1	IP 2
Year 2	Math 4	Physics 3	Circuits 3	Circuit 4
	Signals 1	Telecom	Signals 2	EE for Next Generation …
	Physics 2	IP 3	Elective	
Year 3	minor		Elective	Graduation Project
			Digital 3	
			Signals 3	

Figure 1. Bachelor Electrical Engineering DRAFT

New Master CESE

An interview with Koen Langendoen, Director of Studies

By T. Pouwels

As of next year, the masters Embedded Systems and Computer Engineering will be merged into a new master, Computer and Embedded Systems Engineering (CESE). Starting this year students from both masters follow the same courses together. For this article, we interviewed professor dr. Koen Langendoen. He received his Ph.D. in computer Science from the Universiteit van Amsterdam in 1993 and has been part of the academic staff in Delft since 2001. Now he chairs the Embedded Systems group, and he is the Director of Studies for the masters Embedded Systems and Computer Engineering.

Koen, you are director of studies for the new master programme, can you tell us what that entails?
If only I knew... so previously Nick van der Meijs was Director of Studies for both the bachelor Electrical Engineering and the masters. Even for Nick that was a bit much, so now I have taken over the master CESE. This makes me, to some extent, responsible for the day-to-day operations, and then mostly about the education.

At the moment that entails evaluation of courses, the focus groups and the evasys surveys. Next to that I am also very busy with the accreditation. Once every 6 years programmes must go through this whole process, you must appear before a committee, and I am the one coordinating that.

And this is only for Computer and Embedded Systems Engineering?
Yes, so right now we still have two separate master programmes, from next year on we will have one master programme. We kept the same letters, just changed the order a bit.

Follow up question, how should we pronounce CESE?
I always pronounce it as [kei-si]. ('keesie' for our Dutch readers). It is just too long of a name to pronounce it in full.

Why were the masters merged?
You could also look at this from another perspective, why were they ever split? At first we had, I'm going from the bottom up, Electrical Engineering, then you have Computer Engineering about the hardware and Embedded Systems which has a more specific scope, and then on top you have Computer Science. Meaning they are both in the same area between the other masters. The question whether we should merge them has been going around for a long time. The question then rises what is the difference between the two? For example, if you wanted to you could get both diploma's with the same set of courses.

Some things that had to be considered were the fact that Embedded Systems is originally set up as a 3TU master, together with Eindhoven and Twente. The idea was that this would also lead to a joint degree, but due to differences in regulations this turned out to be more difficult than expected. By now this is a 4TU master, Wageningen also joined, and we let go of the idea that we should keep the same core courses as the other masters.

One of the consequences seems to be that we now have a very large common core. Why is that?
So before Embedded Systems had around 25 ECs of common core, and the other ECs were free to choose. Computer Engineering had some common core, and then some specialisation tracks you also had to choose

What is Computer and Embedded Systems Engineering? *from: tudelft.nl*
In this master you will learn about advanced computing and software systems, system engineering and computer architectures for embedded applications as well as high performance systems. In the programme you will be faced with real-time operating aspects and computer arithmetic. In projects you will experience how all these parts come together when designing an embedded system in practice and an optimised processor. In these projects hardware and software are integrated and you will learn how to design and develop systems which require this multi-disciplinary integration. You can specialize further in the direction of Computer Architecture, Networking, Software or Control.

courses from, so there was some choice but not as much.

Then you can take the intersection and you will end up with 0 courses in the core, or you take the union and you end up with only core courses. So we tried to weed out some courses, but we still increased in the amount of courses compared to the previous programmes. We also introduced the homologation course, which is Software Fundamentals for the Electrical Engineering students and Hardware Fundamentals for the Computer Science students.

How many students follow the master? And what was the target?

Previous years we had around 20-30 Computer Engineering students and 60 – 70 Embedded Systems students, so we aimed for 100 students for the new master. We got a little bit more, which is unfortunate.

How so?

Because it is nice to have a smaller group, it means you have more time for individual students. Some courses were also given for only 20-25 people previous years, so that takes some effort scaling up to 100 students.

Is there already some feedback on the first quarter?

I had a lot of appointments yesterday, so I didn't have the time to look at the evasys results yet because they just came in. However, I also had the student panels at some courses, and they spontaneously also gave some feedback about the other courses so yes, there is a lot of feedback.
Normally there is only the evasys, but the issue is that it's mostly only at the end of a course, so students are mostly still in exam mode and they have forgotten all about the first three or four weeks. And especially now that it's the first year we are doing this, I'd like to have a lot of feedback and then preferably do something with it immediately as well.

Can you shed some light on what we are in for the coming years?

That's a hard question to answer. We are now very busy with refining the first year, so hopefully that will go a lot smoother next year. Then I expect we will be busy for one more year with

Figure 1. Koen Langendoen, Director of Studies for the master CESE.

that at least. I already said we have the accreditation every 6 years, but there is also a midterm evaluation after three years, so that's the next big thing coming up then.

We will have to evaluate the amount of courses in the core, as we said they increased quite a lot. And we also switched to the Rust programming language. We now give three courses in Rust, maybe that should be a bit more in the future. I do believe it's good to also teach some courses in good old C, for backward compatibility and such.

Ê

	1st Quarter	2nd Quarter	3rd Quarter	4th Quarter
Year 1	Advanced Computing Systems	Software Systems	Specialisation	Effective & Responsible Engineering
	Systems Engineering	Real-Time Systems	Embeddes Systems Lab	Specialisation
	Software/Hardware Fundamentals	Specialisation	Computer Arithmetic	Processor Design Project
Year 2	Joint interdisciplinary project, Internship, or set of elective courses	Thesis Project		

Figure 2. The curriculum for the master Computer and Embeddes Systems Engineering.

Stories of EEMCS

From the former head of EEMCS facility management and honory member of the ETV

By J. van der Pol

I was asked to write an article about the EWI-building, also known as the Electrical Engineering Faculty. Or, as people in Delft like to call it: 'Elektro'. On the fifteenth of May 1973 I began working in 'Elektro', and this continued until the fifteenth of October 2005. More than 32 years in the same building, years in which a lot of things happened inside of and with the faculty. Even after these 32 years, when I was working at the department of Real Estate at the TU Delft I still had a lot to do with EWI. It's the only building on campus (and in the Netherlands) that makes my heart beat a little bit faster. Constructed in the sixties and opened in November 1969. For more than fifty years, the building has been a recognizable beacon in Delft.

The Facade

For years we made a giant Christmas tree by selectively turning on and off the lights in the rooms of the high rise. You should know that in those times it was customary for EWI to close earlier in the evenings. On Fridays, the building closed already at 19:00. To make the Christmas tree, we started with a drawing of all the rooms. Then, on the last Friday of the year, together with some students, we went by every room turning the lights on or off as by the drawing. Afterwards we celebrated in the E-kafee, which is now known as the /Pub.

Figure 1. Christmas three on the facade of EEMCS.

EOW

At some point as facility management I got in touch with the Electrotechnische Vereeniging. When organising the EOW (the first-year weekend) we agreed to let the new students sleep in EWI on Friday. That Friday we made the letters EOW on the facade using the same technique as with the Christmas tree. The new students were dropped somewhere in Pijnacker and could walk back to EWI where, as a beacon in the night, the letters EOW towered above Delft.

Abseiling

During the preparations for the OWEE, some representatives of the Delft Student Alpine Club, part of the Delftsch Studenten Corps, asked if during this week they could use the building for abseiling. After discussing the safety of the participants we decided it would happen on Monday. To my surprise however that morning the clock of EWI had been covered and written in chalk was "YETI, wij waren eerder" (we were first). YETI being the Alpine club of K.S.V. Sanctus Virgilius.

Figure 2. Abseiling from the side of EEMCS.

Twice a year the Haaglanden police came to practice abseiling as well. This always happened in the evening. With very fast cars they showed up and quickly descended. The army also came to practice three times. The soldiers had to get up the stairs to the roof every time for the abseiling. Now it just so happened that my brother in law was also in the army and I knew he did the same thing the year before, but at the Euromast. After asking the soldiers why they picked EWI and not the Euromast, it turned out that a general of the navy had said "Rotterdam is een marinestad en de landmacht blijft van mijn toren af" ("Rotterdam is a navy city and the army stays off my tower").

Base jumping

My phone rings. It turns out to be the VARA, a Dutch broadcaster. They have a radio programme called Dakterras, and they wanted to know if it is possible for paratroopers to jump off the building during the show. After discussing safety it turned out it would

be broadcasted during the afternoon. I did have some objections to that: imagine the students and staff seeing people jump down through their windows. Instead they agreed to jump in the evening. After everyone had made their jumps, police showed up. It turned out that base jumping is illegal in the Netherlands.

Police

For two weeks the police was stationed on the 22nd floor of EWI. From here, they looked over Delft and Rotterdam. I was very intrigued to the reason they were observing from here (I had arranged the room, the access and other facilities), but they were not willing to share anything. On one night I heard on the news that the police had done a big raid at a company in Delft. "I'm assuming I will get the keys to the building back quickly now", I said to my wife. Indeed, not two days later this was the case. As it turned out the police had been filming nightly transports that in the end led to a conviction

The ETV XVIII-Lustrum - Tetris

For two weeks, staff and students were unable to look outside during their work. The ETV was very busy laying cables and masking the windows in order to prevent the light in the room to be visible from the outside. On the

Figure 3. The iconic TU Delft letters being lifted to the top of the building.

17th of November (because it had to be dark) the honorary chairman of the ETV, De Kroes, opened the Lustrum-stunt, after which the tetris blocks fell down the entire weekend. That weekend I also celebrated my birthday and from my home the EEMCS building was visible. At a certain moment we ran out of wine and because back then eveningshops did not exist, I quickly drove to the e-kaffee where the ETV board still was. I got a box of white wine and used the ETV phone to call my wife if she could look outside for a minute. There the following sentences were going across the facade: "Jan, sorry we are misusing your building" and "Corrie, the white wine is coming!".

The ETV XIX-Lustrum - SMS for building

This was a fairly similair stunt in regards of the rooms and the light, but also used a different technological aspect. People were able to send a sms message, which was then showed on the facade of the building. With the Lustrumcommittee I made sure that there would be a filter to prevent any rascist or otherwise inappropriate text to appear. Luckily this worked fine and the stunt on this building was again amazing.

Ê

Figure 4. Base jumping from the roof of EEMCS.

The Delfts History of Vlek

A look inside one of the traditions of the Electrotechnische Vereeniging

By M. Vlek and J. van Breukelen

For many years, Vlek has been the tradional drink of the Electrotechnische Vereeniging. The drink is originally from Delft and has a long history. For this article we sat down with Marc Vlek, great-grandson of the founder of Vlek, and writer of the book "Brood, Geloof en Jenever".

How it all began

The Vlek family had been a family of bakers in Leiden for hundreds of years, when around 1850 Jacobus Johannes Vlek (1832-1918) decided to move from Leiden to Delft. There could be several reasons why he moved to Delft. For example, there was a major cholera outbreak in Leiden around that time, but it could also be because of the fact that the population of Delft was relatively Catholic and the Vlek family was Catholic as well. He started as a grocer's servant on the Lange Geer in Delft. In those days it was common to live in the house of your employer, and so he did. Jacobus Johannes was loved by his new family and as a result, he married the grocer's daughter. The family would later get 11 children. The store was taken over in 1855 and it would be run by Vlek for 15 years.

Because he was involved in a large network of Catholic people, he got the opportunity in 1870 to buy the distillery "De Verkeerde Wereld" on the west side of the Brabantse Turfmarkt. The jenever he brewed was immediately renamed Vlek. Later he would also buy the building next to it, called "Nooitgedacht". Meanwhile, "De Verkeerde Wereld" no longer exists, but "Nooitgedacht" can still be found today at Brabantse Turfmarkt 17. On the facade of the building there is still the text "DISTILLEERDERIJ "NOOIT GEDACHT N.V. VLEK and CO". Rumor has it that the distillery used to belong to the famous painter Jan Steen, but this is not true. In fact, it was located on the Oude Delft. In addition, it has also been advertised that Vlek was established in 1550, but this is also not true.

Figure 1. Distillery Nooitgedacht at the Brabantse Turfmarkt 17.

Due to the decreasing demand for jenever, high taxes and fewer exports, the businesses for Vlek started to decline. Consequently, on July 26, 1901, Vlek was declared bankrupt. On March 15, 1902, Vlek was relaunched and taken over by his two sons, Jo and Koos Jr. They introduced innovations that made the production process easier and cheaper.

Often friends would visit the distillery and a lot of Vlek was drunk together. Koos Jr. also invited students from the Technische Hogeschool Delft, which would later become the TU Delft, on a standard day in the week to come by and drink Vlek at the distillery.

Reinier de Graaf Hospital

The Vlek family was also very active in the Catholic emancipation at that time. At the time of the new constitution of Thorbecke in 1848, Catholics were finally allowed to organize and manifest themselves publicly. Catholic emancipation arose around 1860. From that time on, the Maria van Jesse Church was built in Delft and many Catholic (student) associations were founded, such as Virgiel. Catholics in Delft also wanted their own hospital, such as the Oude and Nieuwe Gasthuis, because they provided care for the poor. They started as a retirement home to earn money to take care of the poor. Johannes Jacobus Vlek was chairman of the board and had to make sure there was staff. The way he did this was by going to a nunnery in Dendermonde in Belgium and recruiting nuns there. The Catholic hostipal that was created would eventually merge into the Reinier de Graaf Hospital.

TV Commercial

The Vlek family was commercially very innovative. Just after WWII, they wanted to increase sales of Vlek. Therefore, they approached the editor-in-chief of Elsevier to get a few pages of advertising space for 100 bottles of Vlek. The editor-in-chief thought this was a good idea because he thought it would bring the editors back to the office. So as a result, Vlek got ads in the Elsevier.

Later, TV commercials by Vlek were broadcast on national television. Thus, by the 1970s, the slogan "Even een Vlekje wegwerken!" had become very well-known in the Netherlands. Another phrase that often appeared in Vlek commercials was "Ik heb trek in oude Vlek!". This is a phrase that is remembered by many people.

"Even een Vlekje wegwerken!"

That the Vlek family was commercially very innovative can be seen by the fact that Piet Vlek (1912-2001), grandson of JJ Vlek, went on a business trip to Curaçao in the 1950s, because lots of jenever was drunk per capita in Curaçao. So much in fact that it was almost impossible. The real reason so much Vlek was drunk is that there was a boat market, where people from Venezuela came with fresh vegetables because it was not grown in Curaçao. In exchange, these people received jenever, with which they sailed back to Venezuela. From 1950 to 1967 there was also a soccer tournament in Curaçao called the "Oude Vlek Voetbaltoernooi," sponsored by Vlek.

Vlek leaving Delft

In the years that followed, it was becoming harder for Vlek to keep up. The market for Dutch spirits was declining considerably. Jenever was losing ground to wine, cognac and whiskey, which were being made popular in the Netherlands by large multinationals with huge advertising campaigns. Against these large companies, the family distilleries in the Netherlands could not compete, although Vlek never gave up and even ventured into expensive television advertising. So in 1972, Vlek merged with the company Herman Jansen from Schiedam. A new name was needed. This could not be the (family) name Vlek or Jansen, and so the name UTO was chosen, meaning "Unaniem tot Overeenstemming" (Unanimous To Agreement). As a result, the production of Vlek moved to the jenever city of Schiedam.

The company's brand portfolio changed regularly thereafter, but Vlek remained and became one of Hooghoudt's major competitors in the north of the Netherlands. Whereas at first there were many types of Vlek, such as young jenever, old jenever (Oude Vlek), but also flavours like lemon jenever, eventually only the young jenever remained at UTO, as there were many different brands at UTO. In the mid-1990s, the Vlek family sold the shares and then UTO sold the brand name and production of Vlek to the company Boomsma. UTO meanwhile changed its name back to Herman Jansen in 2011. Nowadays, Vlek is mainly drunk in the North of the Netherlands and in Delft and its surroundings.

The logo

The man depicted on Vlek's current logo is the French playwright Victorien Sardou (1831-1908), who wrote mostly satirical pieces. Why he was chosen is unknown, but they probably somehow related it to jenever.

The old logo is very reminiscent of the VOC logo, as it consists of "V&C," which stands for "Vlek&Co". This was done on purpose because jenever was also exported to West Africa. Ghana belonged to the Netherlands for a time and they traditionally often drank jenever there. Jenever was also often used as currency in those days. Those people regarded the VOC logo as trustworthy because in those days they often got their money right away from the VOC merchants. Therefore, Vlek had a logo that closely resembled the VOC logo.

Connection to the ETV

In the 1980s, Vlek was introduced to

Figure 2. The old label of Vlek with the old logo.

the Electrotechnische Vereeniging by then ETV member Ruud Reichert, a grandson of Jok Vlek (1907-1984). As the only association in Delft, the Electrotechnische Vereenigng has an association drink originally from Delft. Vlek is drunk on a number of official occasions of the association. Prospective first-year students are already introduced to this Delft jenever on the freshmen weekend and after a member's graduation a "Vlekje" is often drunk to celebrate it. Traditionally, the Vlek is poured only by the treasurer, after which the president does a speech. The Vlek is then drunk with the left hand, before you cheer with the words "Good morning!"

Want to know more about the history of the Vlek family? You can read about it in the book "Brood, Geloof en Jenever" by Marc Vlek, available at www.broodgeloofenjenever.nl

Ê

Capture the Flag

Discover Your Inner Tech Genius and Tackle CTF Challenges Head-On

By P. Groet

Capture the Flag (CTF) challenges are a popular activity among engineers and other tech enthusiasts. These challenges involve solving puzzles and hacking tasks in order to capture a flag, which is often a string of text or a piece of code. CTF events can be held in person or online, and they range from beginner-level to advanced.

If you're an electrical engineer or have a background in a related field, you may be interested in participating in the ETV CTF platform (https://ctf.etv.tudelft.nl/). These challenges are a great way to test and enhance your technical skills, as well as to learn new techniques and tools. So, what can you expect at a CTF event? Well, each challenge will typically involve solving a problem or hacking into a system to capture the flag. These challenges can range from simple tasks, such as decoding a message, to more complex ones, such as reverse engineering a piece of software. To succeed in a CTF challenge, you'll need to use a range of technical knowledge and skills. This might include understanding coding languages, file formats, and cryptography. You'll also need to be able to think creatively and solve problems in a logical and systematic way.

"As a teaser for a challenge take the image below"

One of the great things about CTF challenges is that they can be applied to real-world situations. For example, you might be asked to find vulnerabilities in a web application or to develop a new security system. By participating in these challenges, you can gain practical experience and apply your technical knowledge in a hands-on way. So, if you're looking for a fun and challenging way to exercise your technical skills, consider participating in a CTF event. Whether you're a seasoned pro or a beginner, there's likely to be a challenge that suits your skills and interests. And who knows, you might even discover a new passion for puzzling and hacking!

As a teaser for a challenge take the image below. It used to be a readable QR code, but someone spilled his coffee over it. QR codes have redundancy built into them, but this one seems damaged beyond repair. Or is it? Once decoded you can verify it in on the CTF website under the challenge "Coffee break(s)" at https://ctf.etv.tudelft.nl/challenge/38/

Ê

ETV (after) Christmas Puzzle

Send your answer to Maxwell-etv@tudelft.nl and you might win a prize!

This December, Santa has come to the ETV, but unfortunately he forgot most of his presents. He wanted to return home leaving all ETV members empty handed, but as he almost left EEMCS he felt he had a present in the pocket of his jacket. This ADALM 2000 from Analog Devices, he thought, was the perfect present for an electrical engineering student. Only problem was, he had only one. He came up with a way to make the most out of this gift: he left the present with the Maxwell committee and gave them the task to make sure it ends up with the best and brightest student.
With this puzzle you can prove you're this student! Find the answer and send it to us before the 1st of March to have a chance on winning the prize sponsored by Analog Devices and Santa. The answer is a word, it is allowed to use the internet.

	important name: (...) Clerk M.	
	EE is the best (...) of Delft	r a d — 1 right, but in wrong place
	EE, (...) & CS	l c d — 1 right, and in right place
	Telex uses paper(...)	l a r — 1 right, but in wrong place
	British vacuum tube	a r z — 2 right, but in wrong place
	*Not a real word	+1

Answer:

Win the Analog Devices ADALM2000

With a value of €230

Flowchart

Which master program fits you best?

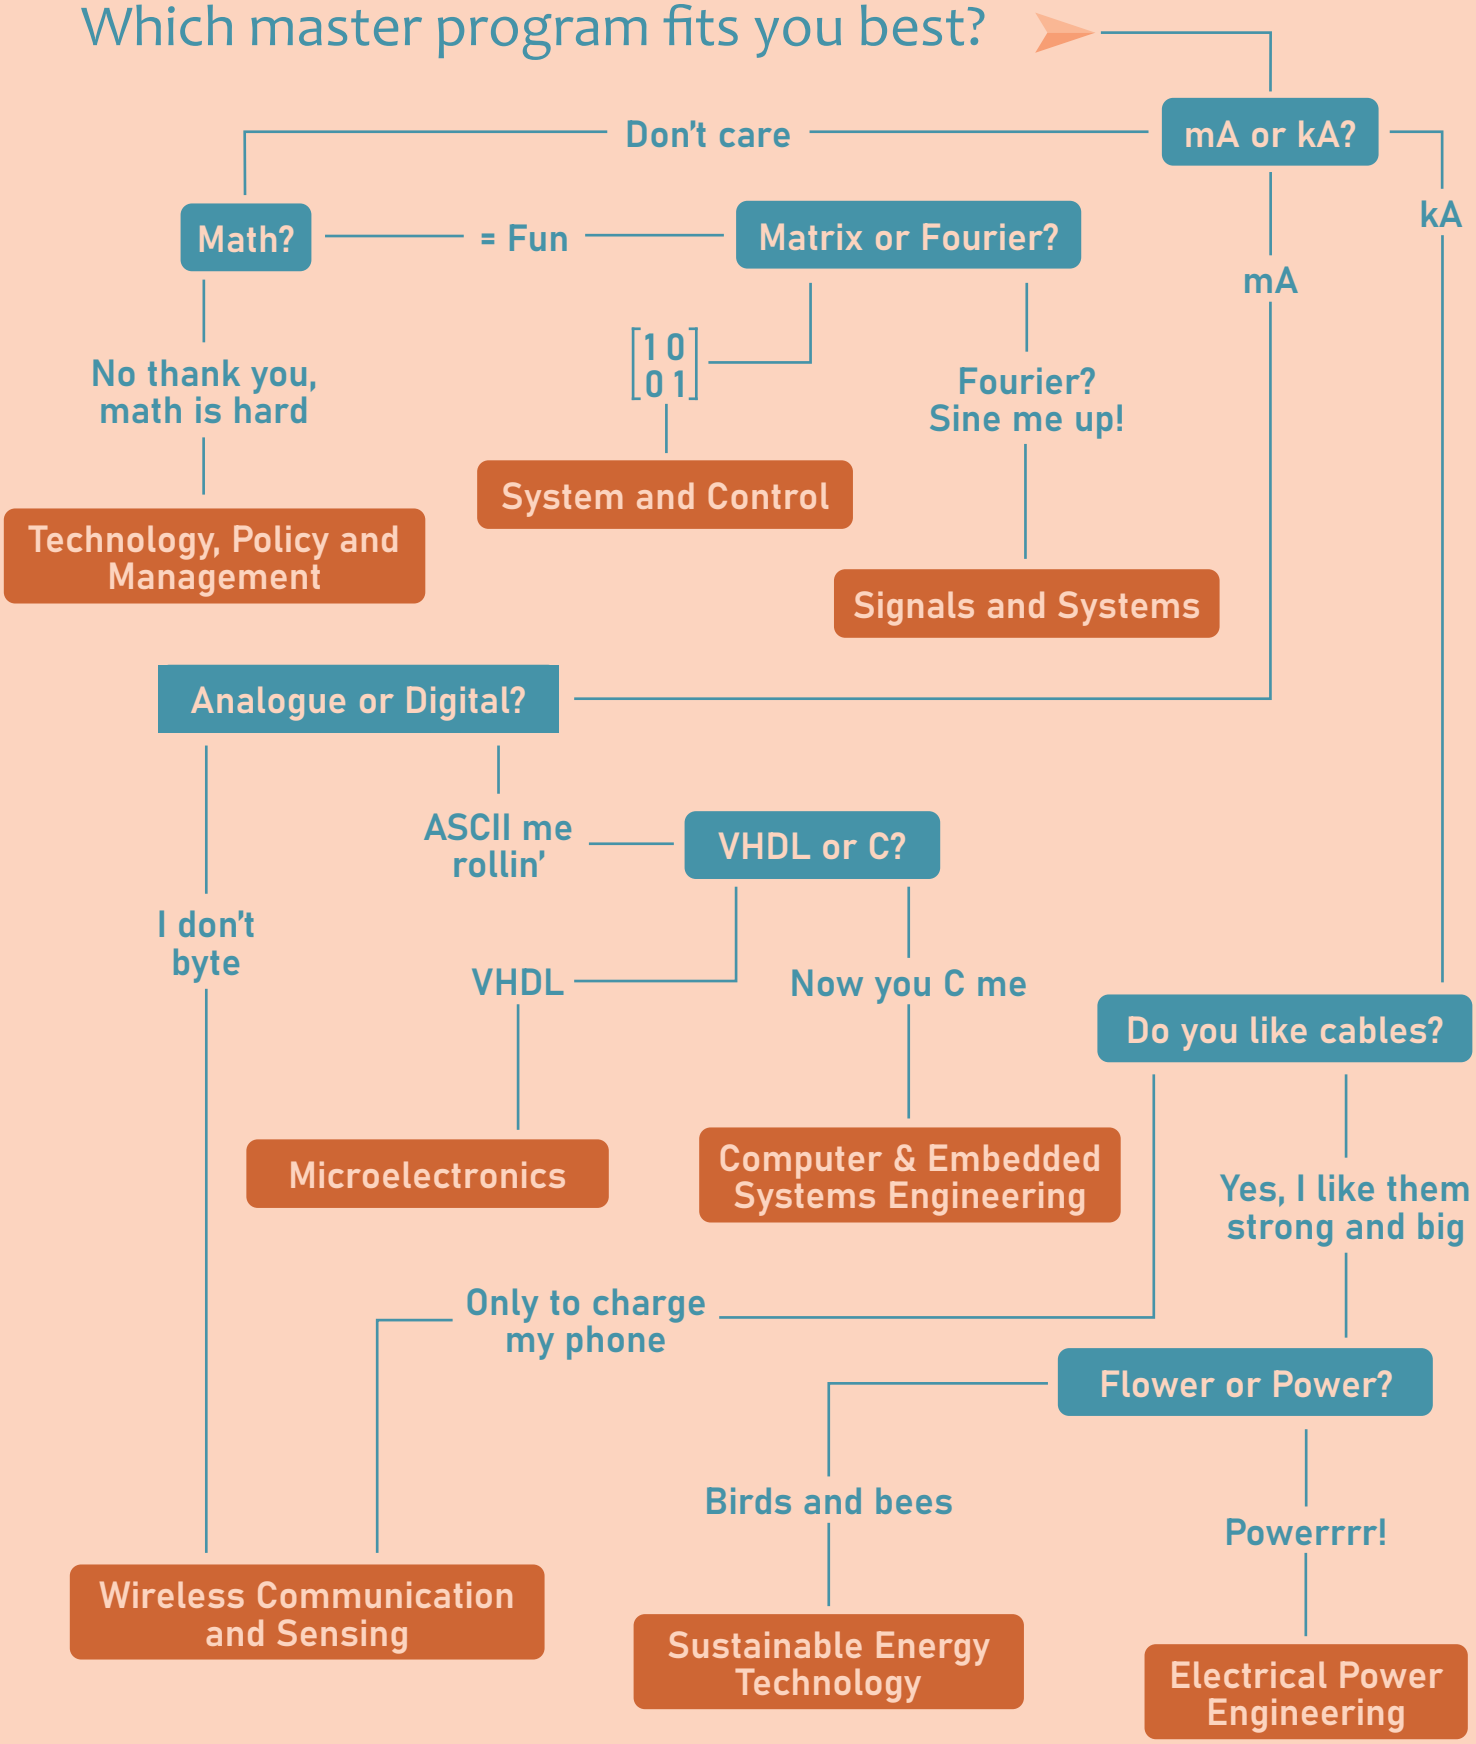

*This is absolutely NOT real advice for choosing a master

/Pub Beer Review

By honorary member prof. dr. ir. A.H.M. Smets

After a long day of studying, many student reward themselves with a nice refreshing beer in EEMCS's very own /Pub. While many student enjoy this moment at the end of a wednesday or thursday, our redection felt there was some untapped potential to make this even more amusing. This is why we sought the help of ETV's honorary member Arno Smets, to help us find which beverages give you the best possible time in the beautiful basement!

Hertog Jan

Draft
Pilsner
€ 1,30
5,0% Alcohol

From the two brands of pilsners on the /Pub menu, the Hertog Jan is my favourite. This has nothing to do with my burgundian mindset and the fact that a mere 20 kilometers from the village Arcen, where the Hertog Jan

"This Limburg pilsner is therefore the one in my refrigerator at home"

brewery is located, I was born. Nope, it is the soft flavour en the pleasant bitter aftertaste in combination with the golden-yellow colour and the full foam head, what makes this beer my favourite pilsner. This limburg pilsner is therefore the one in my refrigerator at home. For student a perfect beer as well, because it can be ordered studentikoos at the bar: "Doe mij maar een Haa-Jeetje".

Duvel

Bottle
Blond
€ 2,20
8.5% Alcohol

This beer only tastes well when served in the original Duvel beerglass (Figure 1). The shape of this glass is just as important as the characteristic Belgian golden colour of the beer. To enjoy this gold, you have to carefully pour the bottom of the bottle in the glass, as to prevent the cloudy particles to stream along with the beer. The Duvel has a strong taste with hints of citrus. This

Figure 1. The characteristic Duvel glass.

all gives the beer a sweet, sour, herbal and bitter flavour. But be warned, it remains a devil!!! On an empty stomach after a busy day, consuming just a few Duvels in the /Pub makes you feel the alcohol quickly like a "duveltje uit een flesje (doosje)".

Jupiler 0.0%

Bottle
Pilsner
€ 1,20
0.0% Alcohol

I was asked to write a review of an alcohol free beer. 'alcoholic free beer' is in its essence an antithesis, like a ham-cheese tosti without the ham and cheese or a vegetarian meatball from the butchery. Nevertheless has this type of beer an important duty to fulfil, it gives the BOB the opportunity to enjoy a glass with golden colour and a foam head while socializing with his or her friends without feeling left out. Another advantage is that 1 glass of

"Nevertheless has this type of beer an important duty to fulfill"

Jupiler 0.0% contains 22% less calories than a regular Jupiler. This saves you a subscription at the gym. How does the Jupiler 0.0% taste? I have no clue, when I visit the /Pub I always go with my bicycle.

Past Activities

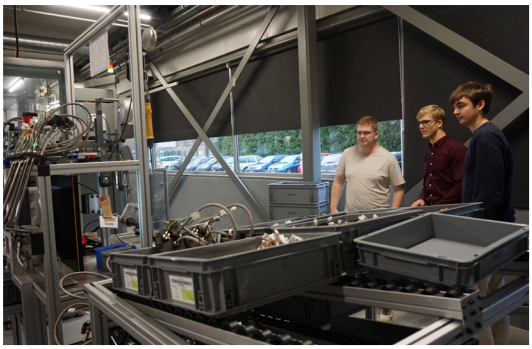

Upcoming Activities

Valentines Lunch
Februari 14

Electrip
March 1-5

Rally
April 1 & 2

EEMCS
RECRUITMENT DAYS
March 20-24

Q3





Dies Natalis
May 8-12

Interfaculty Beerpong Tournament
May 2

Sailing Weekend
June 2-4

Q4

For further information or more activities, visit **etv.tudelft.nl**

Witteveen + Bos

Ruimte voor talent en ambitie

INGENIEURSWERK MENSENWERK

Onze ingenieurs en adviseurs krijgen alle ruimte om het beste uit zichzelf te halen in projecten op het gebied van water, infrastructuur, milieu en bouw. Wat is jouw talent? **www.witteveenbos.nl > werken bij**