De Vonk

Periodical of E.T.S.V. Scintilla

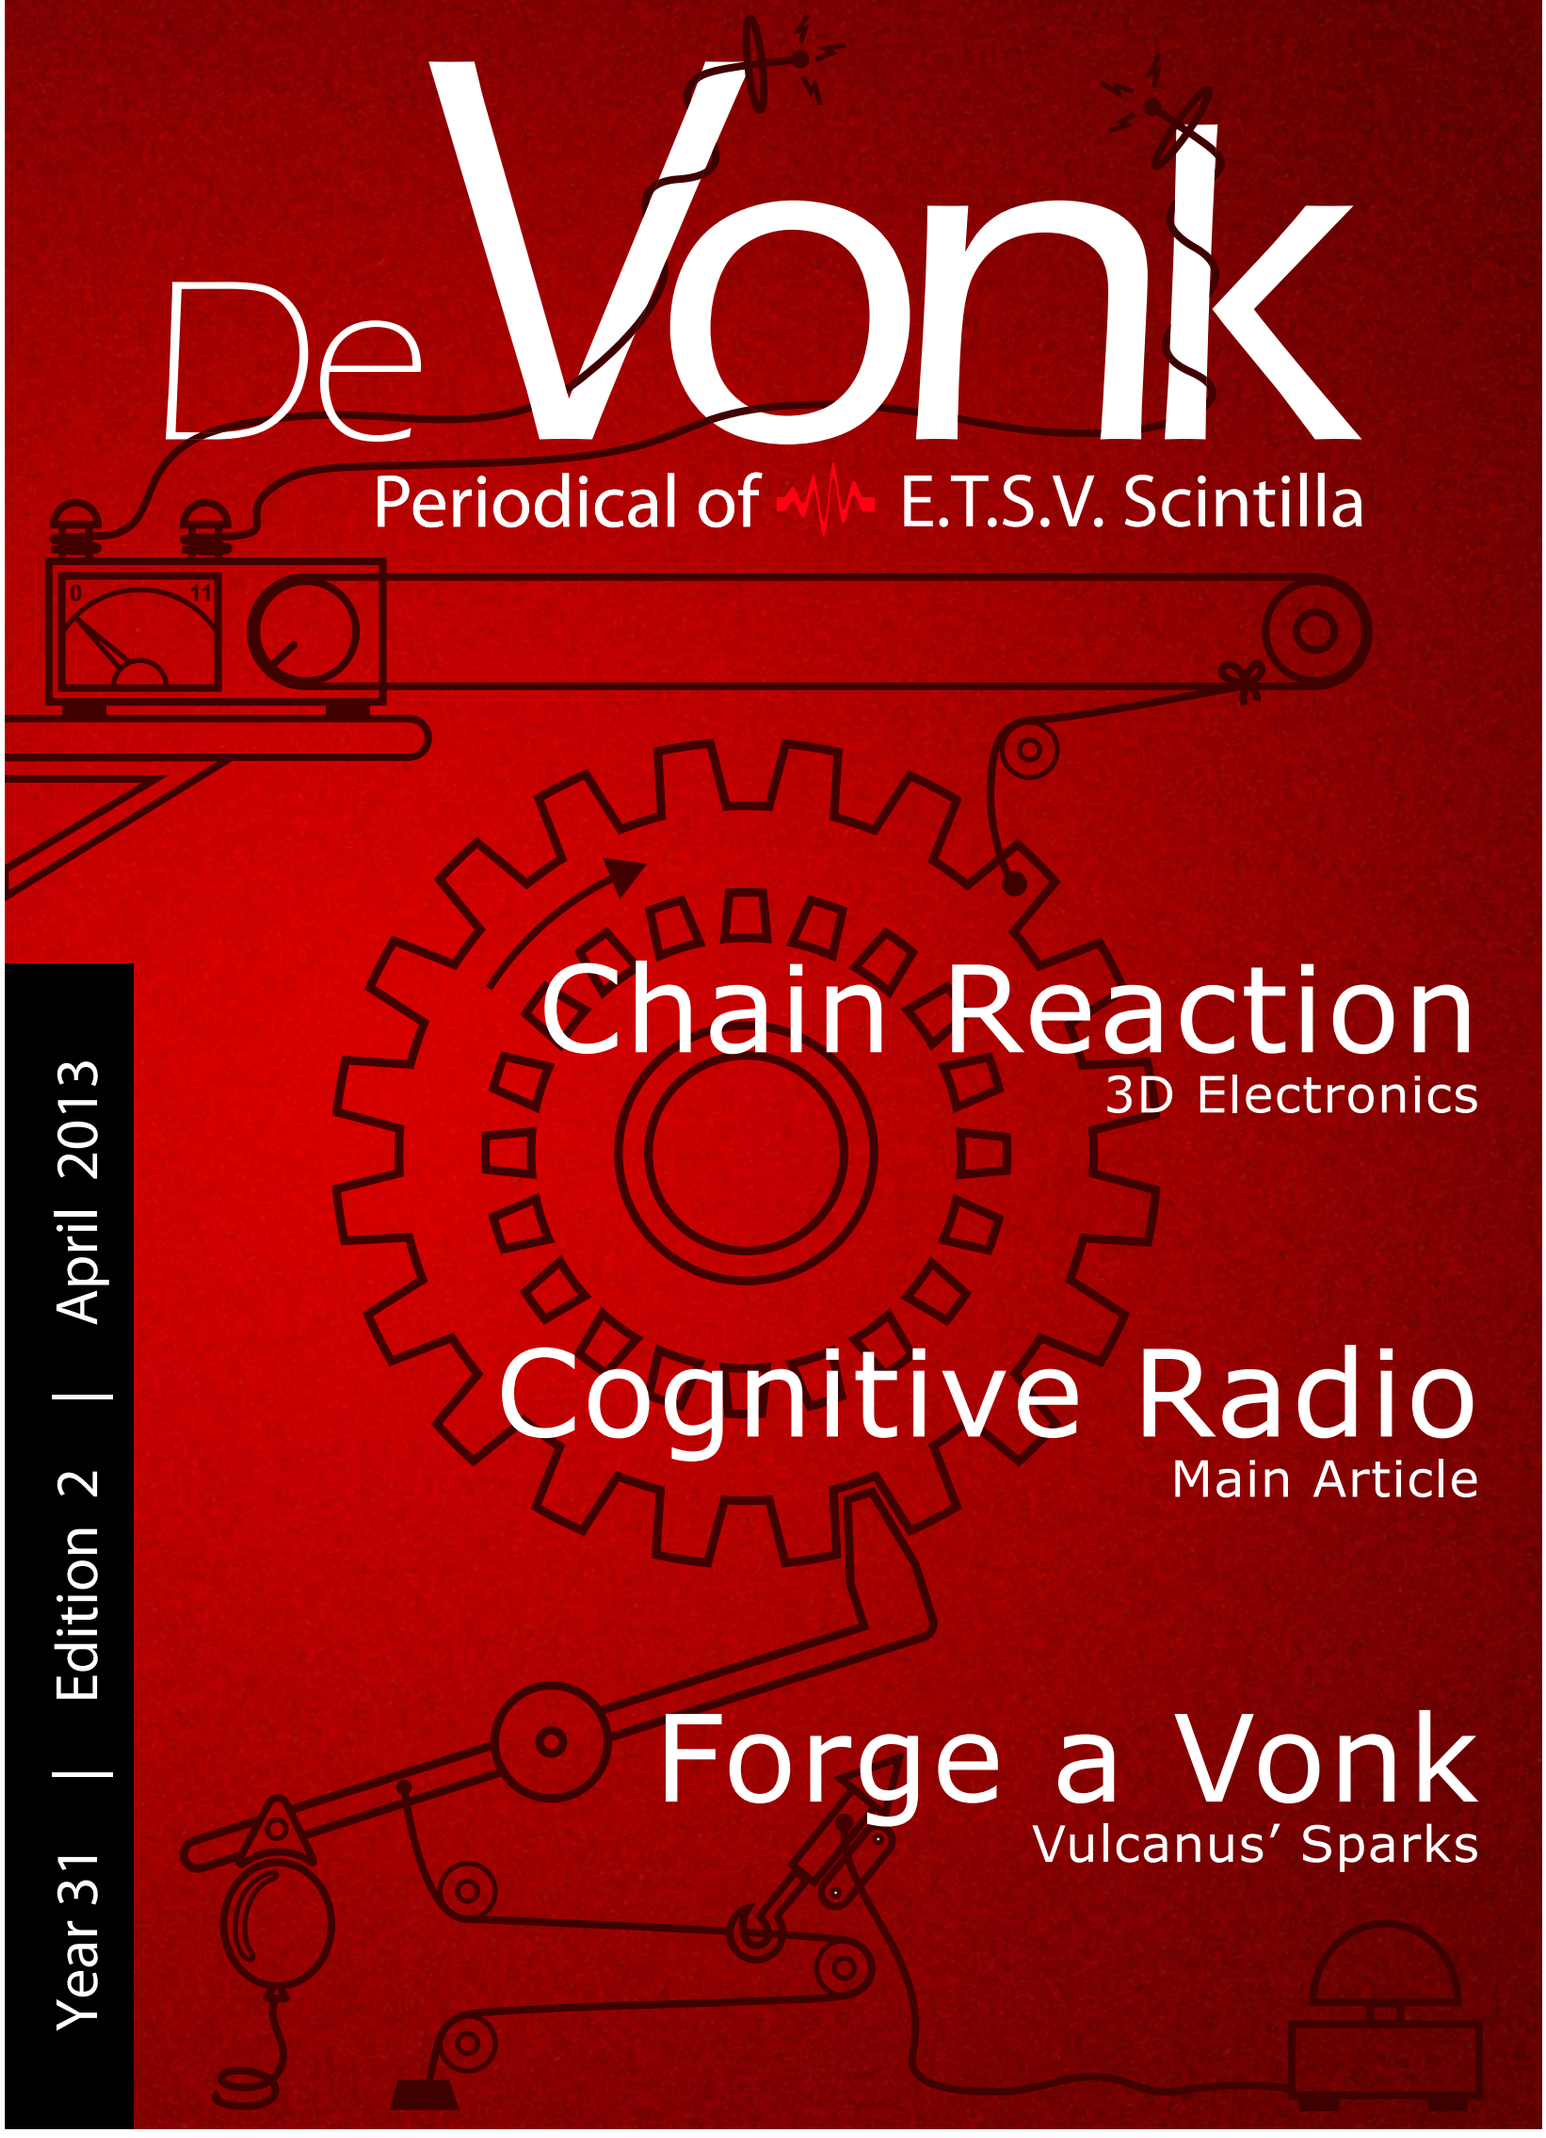

Chain Reaction

3D Electronics

Cognitive Radio

Main Article

Forge a Vonk

Vulcanus' Sparks

Year 31 | Edition 2 | April 2013

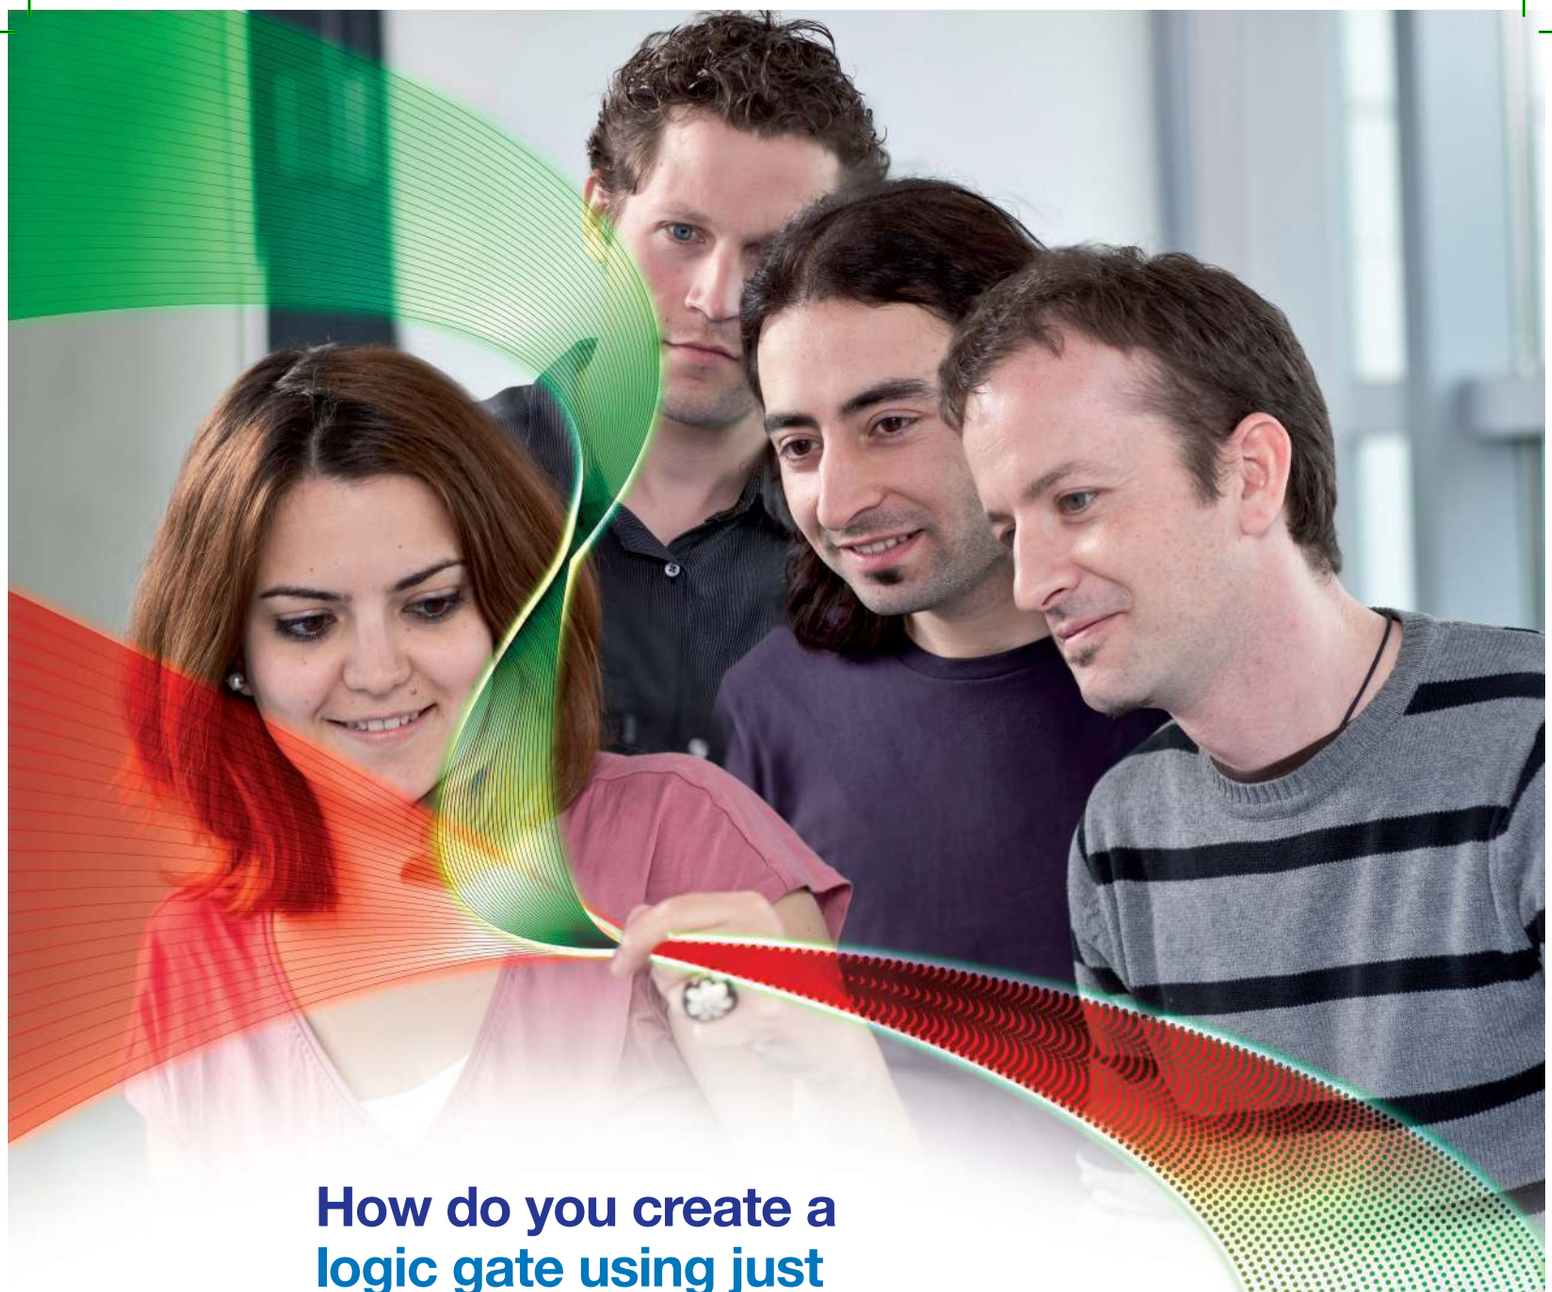

How do you create a logic gate using just 24 silicon atoms?

Join ASML as an Electronics Engineer and help push the boundaries of technology.

As one of the world's leading providers of lithography systems for the semiconductor industry, we bring together the most creative minds in science and technology to help produce increasingly cheaper, faster and more energy-efficient microchips.

For the past 25 years, we've been helping to realize Moore's law. Now we want to go even further: to enable the tripling, or even the quadrupling of chip-feature density every two years. That's why we need talented Electronics Engineers. People who can, for example, increase the speed and the precision of our systems integration, and thereby enable future logic gates no bigger than a few silicon atoms. It's just one of the many electronics challenges at ASML. And there are countless more – all targeted at achieving more than Moore's prediction.

If you're up for these challenges, we'll put you in a multidisciplinary team and give you plenty of freedom to experiment and learn new skills. What Moore could you want?

www.asml.com/careers

ASML

For students who think ahead

Halfway

Author: Koen Zandberg

At the time of writing, we're about halfway through this college year. The new quartile has begun and most of you will probably be enjoying your new courses (or trying to pass some again). When you're reading this, you'll probably know whether you passed them or not. As a board, it is about we showed some of the things we've accomplished and tell you about other things we failed to do. We've been busy for half a year, and we've had enough time to get some things done.

One of the new things we've set up is called Satis. Most of the work on this project is done by our (beloved) webteam. Satis is a web environment that should help most committees with their administration. Every committee has meetings, makes minutes of these and some members even have tasks they should do. Satis is supposed to help with this. As of now, Satis has only a few modules. A task module to keep track of the tasks of members, a module to reserve a room for meetings and a failure module to keep track of failed tasks. At the moment there is even a module which enables committees to upload their own promotional posters to the television slide show. The webteam is working on more modules, so when you're reading this, they probably have finished some more. Satis is fully in development and the webteam need all the feedback they can get. Try Satis out [1] and give them your feedback so they can make it even more awesome.

Something most committee members probably don't know anything about are the so called "commissieinstellingsbesluiten" or committee rules and regulations in my best English. This is a document that every committee has. In this the rules and regulations of the committee are described. Most of these documents are filled with what a committee should be doing with their time. Unfortunately, most of these documents are horribly outdated and as far as I know, most active members do not even know that these exist or haven't seen them in years. For example, according to their "CIB", the Borrel should organise drinks every other week in the Tombe. As far as I know, the Tombe doesn't exist anymore. Thus for now, our plan is to update these CIB's and make them more accessible to our members so everyone knows what they should be up to.

Well, I have used most of the space I get in this corner of the Vonk. Nothing more to report here, so read on and enjoy this edition of the Vonk.

Op de koningin, op Scintilla

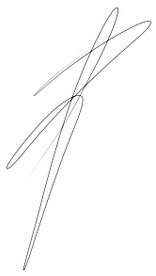

Koen Zandberg
President of E.T.S.V. Scintilla

[1] http://satis.scintilla.utwente.nl (only on intranet)

Hardlopers zijn doodlopers
Thursday April 25th, 2013
16:00h, Sportsbar Abscint

41st Batavierenrace
Friday, April 26th, 2013
15:00h, Nijmegen - Enschede

77th Cantus Scintillae
Thursday, May 2th, 2013
20:00h, Abscint

Symposium
Wednesday, May 15th, 2013
9:00h, Spiegel 1

Masthead

De Vonk

Periodical of E.T.S.V. Scintilla.
Published four times a year in the amount of 630 copies.

Year 31, edition 2
Februari 2013

Editorial team

Tim Broenink, Erwin Bronkhorst, Arno Geurts, Tijmen Hageman, Fieke Hillerström, Maikel Huiskamp, Peter Oostewechel, Ray Tanuhardja, Rowan de Vries, Marcel Wenting.

Print

Gildeprint, Enschede

Editorial office

E.T.S.V. Scintilla, University of Twente,
Postbus 217, 7500 AE Enschede,
0031 53 489 2810
0031 53 489 1068
vonk@scintilla.utwente.nl

Material

vonkkopij@scintilla.utwente.nl

All members of Scintilla receive De Vonk free of charge by post.

ISSN 0925-5421

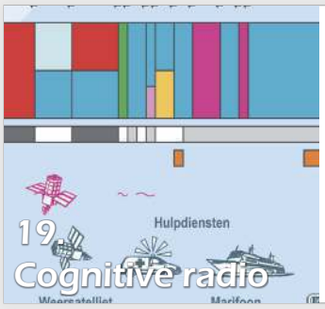

Due to the high amount of wireless transmitters in our society, the available and free frequencies in the air become harder to find and use. However, the demand for available bandwidth is increasing with each new wireless technology that is introduced. To find and use the scarce free frequency bands, cognitive radio (or CR) can be used. In this Dutch article, you can read what research about cognitive radio is done at the UT.

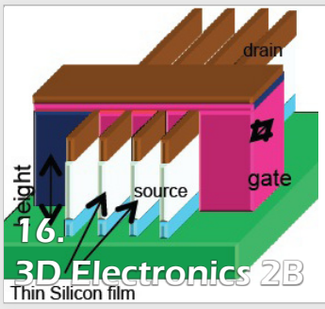

In this Vonk, you will find a new, recurring subject: the Chain Reaction. The goal of this article is to give someone else the opportunity to react on the previous main article. In this edition, some comments on 3D electronics are given, based on the main article of the previous Vonk about 3D self-assembled electronics. Can the self-assembled electronics from the previous article be used with analog, digital and mixed signal electronics?

Dirk-Jan developed a new technology in analog to digital conversion, by making use of a new time-based method. The method analog to two ways of measuring the speed of a runner: you can measure the amount of time that the runner needs to run 400 meter, or you can measure how far the runner can run in one minute time. Each method has its own benefits and drawbacks. The same holds for the ADC domain and in this article, the developed technique is being discussed.

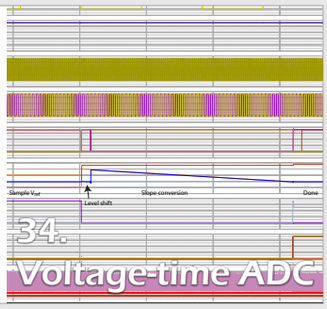

Traditionally, Scintilla organizes a symposium every year. This year, the symposium will take place in May and the subject of the symposium is "Empowering the future". In this article, the organizing committee will explain what you can expect from this day and what kind of speakers will be there to show you their view on this subject. If you still don't know if you want to attend, read the article and enrol for this event.

Editorial

Connectivity

We are always connected in some way to something. Almost all people are connected with each other via Facebook, whatsapp, twitter or some other way with their smartphones. Cities are connected with roads and railroads, and sending a message between these cities can be done in seconds via mail. We are all so depended on these connections to work, or just be there, that we do not know what to do when they are gone. For some people it is the end of the world when they have no reception on their phone. What if they miss something in those minutes their phone has no connection.

But sometimes you make new connections when the ones you trust on are gone. Take for example when the trains are stopped because of a power outage. Normally when you sit in the train you look out the window or read a book, but there is no interaction with the other people in the train. But when the trains are stopped everybody is talking with one another. Some people start complaining to each other that the trains are not running and the staff can't give you directions, some people try to find their bus to the next station while others start joking about everything they can think off. All these people are now connected in some way. Well until the bus they are waiting for arrives, than it is free for all and the new connection disappears as fast as it appeared.

Maikel Huiskamp

News for the Electrical Engineer

Author: Tijmen Hageman

MIT demonstrates fastest p-type transistor

Researchers at the Massachusetts Institute of Technology have developed a record setting p-type transistor. The device is twice as fast as previous experimental p-type transistors and almost four times as fast as the best commercial p-type transistors. Rather than silicon, the tri-gate device is made of the semiconductor material germanium. The increased carrier mobility results from a special fabrication process placing the germanium atoms closer together than they would find comfortable. This `straining' process is performed by growing the germanium on top of several different layers of silicon and a silicon-germanium composite. The germanium atoms naturally try to line up with the atoms of the layers beneath them, compressing them together.

Source: web.mit.edu

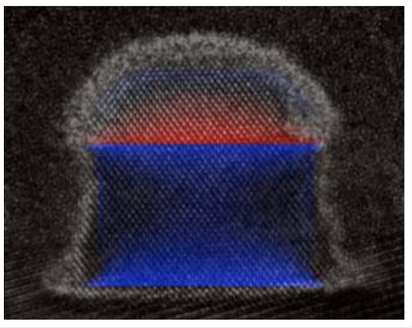

The experimental transistor, blue indicating the areas of `strain', where atoms are placed closer than normal.

IBM shows breakthrough in nanophotonics

For over a decade, IBM has researched the integration of photonic components into electrical circuits, referred to as `silicon nanophotonics'. Recently, they have announced a working prototype created with 90 nm semiconductor technology. Photonics have shown that they allow much faster communication compared to conventional electronic methods. The integration of photonic and electrical circuits in a single chip provides a cost-effective solution for high-speed data communication. The prototype of IBM integrates optical components such as wavelength division multiplexers (WDM), modulators and detectors alongside CMOS electronics, achieving over 25 Gbps per channel.

Sources:
www.tweakers.net, www-03.ibm.com

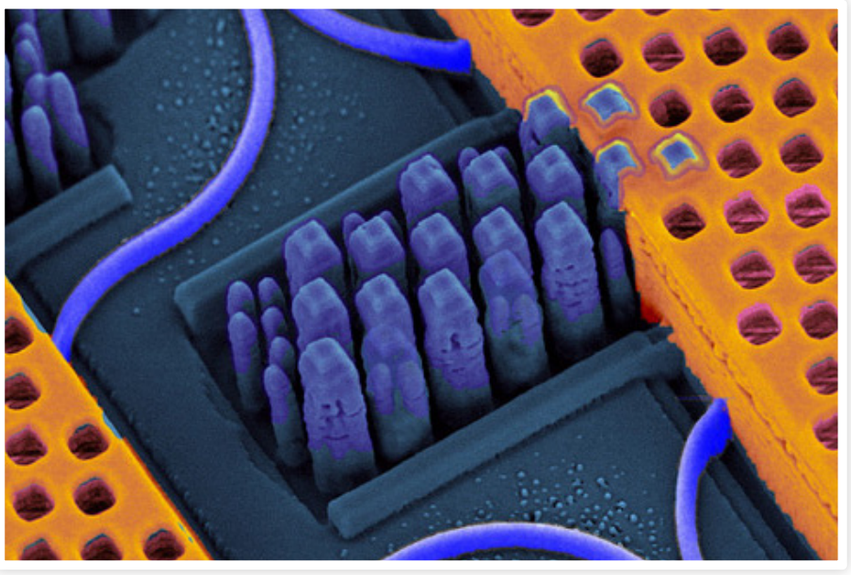

Self-repairing flash memory developed

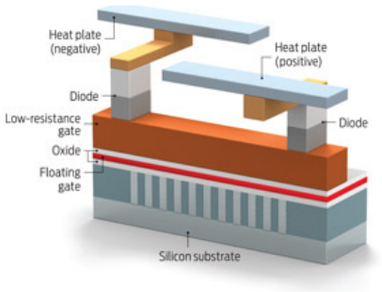

The company of Macronix has developed a self-repairing flash memory unit. Although flash has many advantages, it only supports up to roughly ten thousand read/write operations. This is caused by the fact that flash cells make use of a `floating gate', a layer of material embedded within layers of insulation. By inserting or removing electrons from this layer, a bit is written. Mass read/write operations degrade the insulation, however, and the cell will eventually fail. Macronix figured that when heating the memory cells to 800 degrees centigrade for several milliseconds, the cells can be renewed. They found that the memories held well over 100 million cycles using the technology. The technique is not new: earlier a similar result was obtained by heating the entire chip to 250 degrees centigrade for several hours. This research yielded another surprise: it was found out that heating allows for faster erasing, a process that was thought to be independent of temperature.

Source: spectrum.ieee.org

Toshiba develops low-power MRAM

Toshiba has announced a new version of the prototype memory element referred to as `spin transfer torque magnetoresistive random access memory' (STT-MRAM). MRAM is a next-generation memory technology, exploiting the resistive properties of magnets in different orientations. MRAM is a promising all-round memory technology, as it combines speed, scalability and non-volatile. The prototype of Toshiba has reduced the power consumption by a factor 10, making its power consumption lower than that of SRAM. Especially the fact that MRAM is non-volatile makes it a suitable alternative, also retaining the information in powerless situations.

Source: www.techpowerup.com

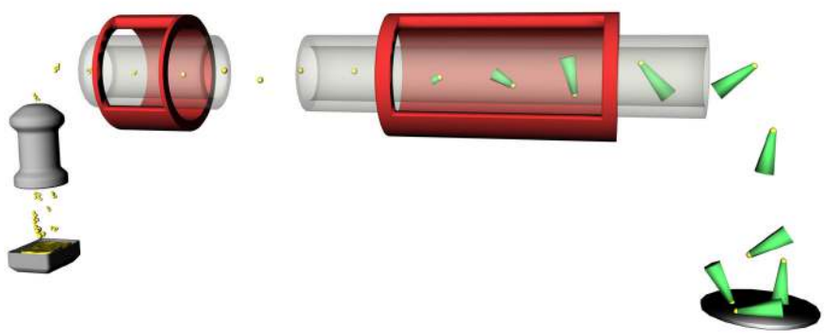

New production process semiconductors developed

Researchers at the Lund University of Sweden have developed a new method of creating semiconductor electronics. Instead of starting from a silicon wafer, the researchers have made it possible for the structures to grow from freely suspended nanoparticles of gold in a flowing gas. A process of self-assembly grows the semiconductor material around the gold particles. The growth can be controlled by temperature, time and the size of the gold nanoparticles. With the help of a series of prototype ovens, the researchers expect to be able to develop several types of nanowires, such as pn-diodes. Not only is the production method very fast and cheap, it also provides a continuous process instead of the traditional batch process that is currently the norm in wafer fabrication.

Bron: www.lunduniversity.lu.se

Graphene/nanotube hybrids developed

Researchers at the Rice University have developed a very high surface density electrode. For this purpose, they first grew a graphene substrate (sheet of carbon, 1 atom thick) on metal. Then, they grew a forest of carbon nanotubes from this substrate. By growing them from the substrate, the electrical contact between the nanotubes and the metal electrode is seamless. The resulting electrode consists of a forest of nanotubes with a height of 120 microns, delivering more than two thousand square meters per gram of material. As the quality of an electrode increases with increased surface, this is a remarkable number.

Source: news.rice.edu

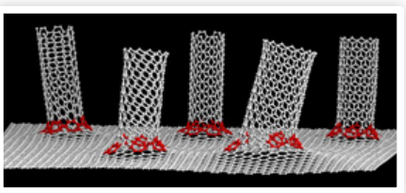

Educational update

Author: Laurie Overbeek

The big educational changes that have been introduced in 2012 also have their influence on the education in 2013 and beyond. I will highlight the biggest changes that were announced recently and will become active within the coming months. Still, a lot is unsure, so watch your e-mail for messages from the university or read the updates that Scintilla will provide.

Deadlines for Bachelor and Master assignments

Bachelor

Since September 2012 it is mandatory to set a deadline for your bachelors assignment. When you have the first appointment to discus your bachelor assignment with your supervisor you should set a start and end date. The performance in the assignment shall be assessed on the closing date, regardless of the stage of your work. If the grade for the assignment is below a pass grade, the chair may give the student the opportunity to continue working on the assignment so as to meet the requirements for a pass grade. The extra time allowed shall however be limited to 3 credits. The grade for the assignment may not exceed a 6 in this case. If the student's performance is still unsatisfactory after the extra time has run out, the student has to do a new assignment on another theme and with a different supervisory committee or under the authority of another chair.

Master

For the master theses a deadline has to be set when the planning is discussed. When this deadline is reached the work will be reviewed and graded. If the grade does not meet the requirements for a pass grade the student gets two months extra. The grade for the assignment may not exceed a 6 in this case, if the grade is still too low a new master assignment must be chosen.

Rescheduled resits 4th quarter

This year the resits of the 4th quarter will not be held in the last two weeks of the summer holiday. Instead these resits will take place in week 30 (22 July -26 July). The regular courses of the fourth quarter will be tested in week 26. The resits of the third quarter will take place in week 27.

Clustering bachelors degree programs

The faculties will disappear and instead all bachelor programs will be divided into clusters. Electrical engineering will be in the "information & Communication Technology" cluster . The clusters are made in such a way that the programs within the clusters can share modules. This does not mean that a program cannot share a module with an other cluster. Instead of a program director for every program we will now have seven directors of education, one for every cluster. The director of education will lead a cluster of courses and mainly fulfill management tasks. They will be responsible for the quality and cohesion of courses within their clusters and pursue new teaching methods.

> "The faculties will disappear and instead all bachelor programs will be divided into clusters."

Why a clustering of the educational programs?

The UT is currently working towards the redesign of all bacholar programs according to the Twente Educational Model (TEM). This model is based on modular educational projects. The University Board wants to group the educational programs in seven clusters. Educational programs with a substantial overlap should benefit from the cluster environment as it simplifies the design process and it should also encourage the educational programs to share modules.

SolarTom

Author: Tom Kooyman

Last time I wrote that the Solarteam project is very similar to other projects. It isn't anymore. It becomes more and more clear that we do not only have to design, but actually build a solar car. From scratch. This means not only assembling every piece of the car, but also finding all of the parts. For free, if possible, and otherwise we also have to collect funding to pay for this. This means that a lot of time is spent looking for companies that are prepared to sponsor parts. Sometimes small (two heatsinks, a couple of brackets, a few connectors) and sometimes large (6 m^2 of solar cells, a complete battery-pack, a few motors).

Another big difference is that the standard scientific layout (describe current situation, propose improvement, test, compare) can't be used on everything that has to be designed, mostly because there is a lack of time. The word multitasking got a new meaning for me. A common day at the moment consists of calling and e-mailing multiple companies, soldering a small PCB, looking into software somebody else wrote, doing some measurements on solar cells and a SolidWorks layout, and finishing the day by preparing a meeting.

On the more technical side of the story I can't bring you a lot of news right now. It is very exciting for us to slowly see the first parts coming in, but I can not show them to you yet. The frame of our mockup arrived last week and by the time you are reading this it should be fully functional. A large box in the mail from Vermeulen Print Service brought us a large panel with all of our designed prints, so we could start soldering and programming.

> "We are looking for the best even money can't buy yet."

In the meanwhile, we are busy with our search for solar- and battery cells. It might not seem to be much of a challenge to buy these components, but it is a challenge to buy state-of-the-art cells. We are not looking for the best money can buy, we are looking for the best even money can't buy yet.

The same is true for the electrical motor that will drive our solar car. In order to know what the motor is capable of, it has to be tested on a motor test bench. And again, I had to find one myself, for free if possible. There is a small company in the south of Holland, called Dynostar, that could help us out. Their company designs and build dynamometers, and they happened to have an old test bench lying around. Have you ever had a scooter? And have you ever had it tested by the police? Chances are that same ROLLERBANK is now used to test our motor.

Right now, we are at the end of the design stage for most parts, but we are not even close to having a functional car yet. After the first tests we will know how much redesigning has to be done and I will keep you up to date in the next edition of the Vonk. If you can't wait that long, there is a blog post every now and then at *www.solarteam.nl*.

Wireless for €5,-

Author: Tim Broenink

Wireless communication is one of the big steps towards making great projects. For what is more awesome than a small robot driving around your room? Five of these robots communication wirelessly with one another. This brings me to this hobby project; to try to set up a wireless connection using a NRF24L01 module, which can be found at the STORES for just € 3,85.

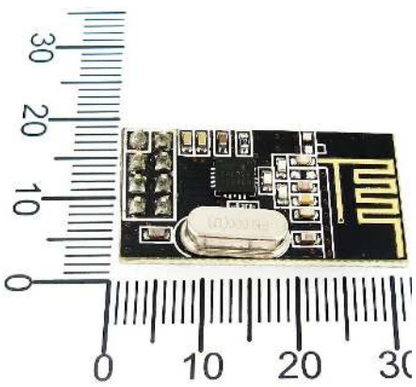

Figure 1: The NRF24L01 Module

The first thing that can be noticed when opening the datasheet of this module is that it contains the things you would expect on a wireless module, including modulators, filters and amplifiers. One of the more unique properties of this little guy is the baseband processing engine. It contains system for the transmission, validation, acknowledgement and retransmission of 'Shockburst' packages. This means that the transmitter handles everything needed for getting your data to the other side. You only have to provide the data.

If you want to use this module there is one thing you should keep in mind. It has a maximum supply voltage of 3.6 V. This shouldn't be a problem for most microcontrollers, because they can operate at 3.6 V or 3.3 V, but it is something that can go wrong.

The pinout scheme of the module can be seen in figure 2. The VCC an GND need no explanation, neither do the standard SPI pins (nCS, MOSI, MISO, and SCK). The only two marginally special pins are the CE (chip enable) and the IrQ (interrupt request). CE will make your transmitter actually send or receive data and IrQ signifies that something has happened in your chip. That was the simple part. We can now connect the module to our project. For this purpose I used an Atmega88. I connected the Power and SPI pins to their equivalent on the microcontroller and connected CE to PB1, nCS to PB0 and IrQ to PD7.

> "Congratulations! You've just send data over the air."

We can now talk to the module. If we look in the datasheet we can find that the module expects a 8-bit command when nCS is pulled low. At the same time it will output its status register. After this it will listen for or transmit up to 5 bytes of data, depending on the given command. To test whether the module is responding we can now send the NOP (0xFF) command. This will only return the status register. The status register should be (0x0e) on a newly started transceiver.

Sending commands is one thing, but we need to do a lot more if we want to actually transmit data. If we look at the state diagram from the datasheet (figure 3) we can see that in order to transmit the data, we will first need to switch to standby mode. After that we can set the data to transmit and transmit. So we will first need to set the PWR_UP bit in the configuration register. In order to neatly write this data we will first need to read this register, change the bit and then write it back to the module. Now the module will actually do something. Lets take a look at the other settings of the transceiver. The default transmit settings can be found in the table of figure 4. For now we will leave most of these settings be. One thing we will change is the RF channel, for the 2.4 Ghz band is also used for Wi-Fi. If we look at the current standards for Wi-Fi (IEEE 802.11n) we can find that the its band stops at 2.4835 GHz (see figure 5). For this purpose we will put our RF channel at 2.496 Ghz. We can set this frequency using the RF_CH register. According to the provided datasheet the RF channel will be 2.4 + (0.001 * RF_CH) GHz. So we will set it to 96.

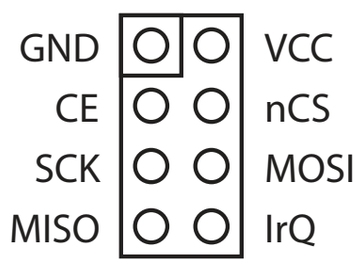

Figure 2: Pinout of the module

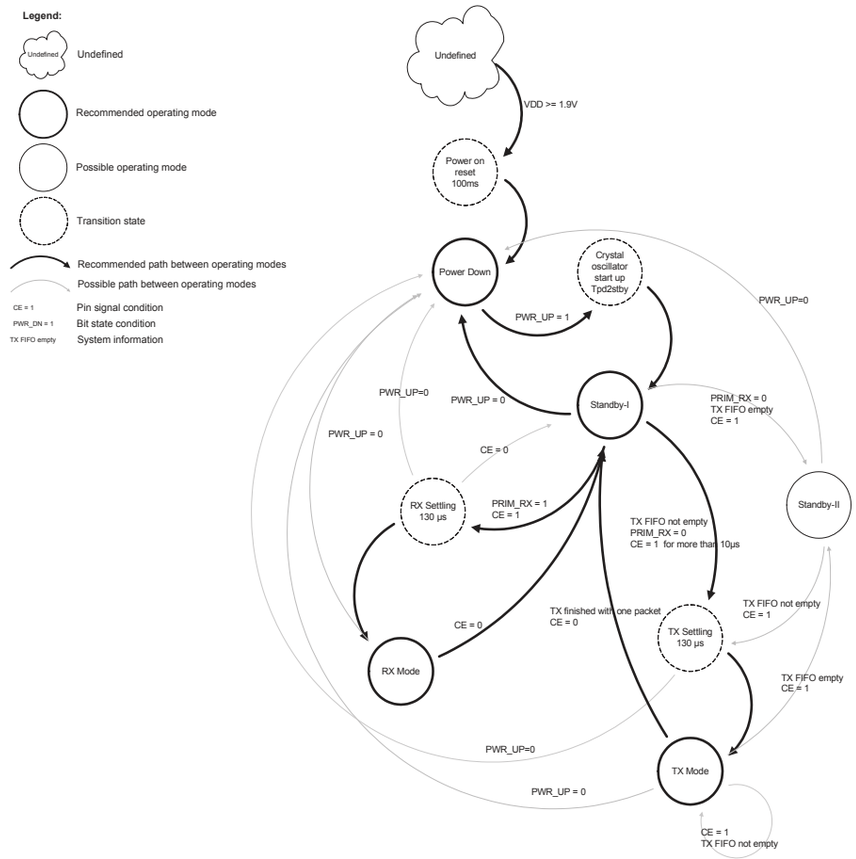

Figure 3: State diagram from the datasheet

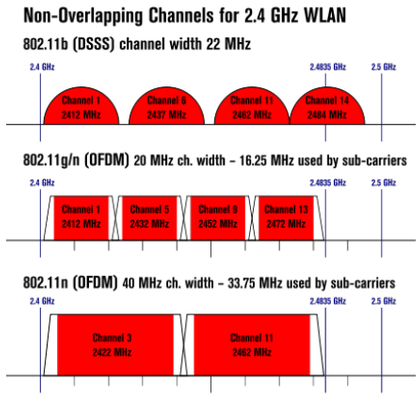

Figure 5: Wi-Fi Channels. As we can see in the diagram, the newer Wi-Fi standards use the channel up to 4.835 GHz. All standards remain below 2.5 GHz.

In order to actually transfer data, we will have to fill the TX buffer. We can do this by sending a W_TX_PAYLOAD command.

We can now write data to the TX buffer. In order to check whether this data has arrived in the module we can check the TX_EMPTY flag in the FIFO_STATUS register. If we then pull CE high, the module will start to transmit. After some time we can read the status again. You will notice that the MAX_RT bit was high. This is caused by the fact that we don't have a receiver to ACK the packet. So our transmitter will retransmit the package 3 times and then give up. In order to fix this we will need a receiver. The receiver has a few different settings, but not many. We will need to set the PRIM_RX bit in the CONFIG register and we will have to set 'receive pipe 0', which has the same address as the transmitter, to a width of one byte. This is done by setting RX_PW_P0 to 1.

If we now send a byte using the W_TX_PAYLOAD command and pulling CE high on the transmitter.

If the receiver is enabled during this transmit, CE high, PWR_UP, then the transmitter should receive an ACK packet from the receiver and will change its status to 'Data sent' (TX_DS). We can now recover the sent byte on the receiver side by using the R_RX_PAYLOAD command. This should return your sent byte. If so: congratulations! You've just send data over the air. If not, you should try to debug using the datasheet and fix it.

Now that the ether is yours to use, you can hopefully think of some awesome projects to make with this. I know that I do and you might see those future projects in another Vonk.

Setting	Default Value
Address Width	5 bytes
Number of Retransmissions	3 tries
RF channel	2.402Mhz
Data Rate	2Mb/s
TX address	0xE7E7E7E7E7

Figure 4: Default transmitter settings

TE Connectivity

Author: Tim Broenink

As Commissioner external affairs of Scintilla I get to visit a lot of different companies at different locations. Once there I often get a short (or not so short) tour of the facilities. During these tours I can see some of the stuff a company is working on. One of my more memorable visits was my visit to TE connectivity.

As far as my company visits go, this one started out quite well. TE connectivity is situated on a fairly large terrain, surrounded by other companies. I could however easily spot their building. TE connectivity is a company corking in many different sectors: Transport technology, Network technologies, Industrial and fabrication technology and consumer goods. Their facility in Den Bosch however specializes in Fiber, Connector and Cable assembly. This was also the subject of my tour. I got to see everything that was involved in the designing, testing and producing of these assemblies.
The first part of my tour was to a large lab filled with large machines. The purpose of these machines remains a bit of a mystery to me, but some of it was explained. The purpose of most of the machines was the same, to manipulate glass fibers in such a way that they could be used. This started simple, with a laser cutter designed to precisely cut a fiber such that it could be inserted into a connector. However it quickly got more complex. One of the most impressive machines I saw there was a large 3-d printer like structure. This monster of a machine was used to "print" glass fibers into a plastic substrate. Resulting in a network of fibers locked in plastic. This could be integrated in circuit boards in order to provide a middle ground between on-chip waveguides and loose fibers.

> "This monster of a machine was used to 'print' glass fibers"

After this large lab, we came to a more familiar environment, a lab like we would recognize here. In this lab I saw a lot of different, mostly electrical tests. From simple current and voltage ratings, to complete measurements of the response of a connector. The setups included maximum current tests. Which used very thick cables and a setup shielded from people with Plexiglas. As you can imagine what happens when you exceed a maximum current for these cables, you can probably understand while this shielding was necessary. But these are the tests you expect on a electrical connector, however there were a few more unexpected tests as well. It is surprising how many aspects of a cable, connetor or assembly need to be tested in order to determine its suitability.

A complete section of the lab was dedicated to durability test. I saw their environmental chambers which cycled the connectors through different temperatures and humidity levels. It went from freezing temperatures to approximately a hundred degrees Celsius. Combined with a humidity ranging from almost zero to a hundred percent, tis put a lot of stress on the connectors. I saw dust chambers which tested the connector in dry and dusty environments. It even had a little wiper on the front, which was used to clear the glass window of dust so one

could see the actual test. I even saw a gorilla test setup, where a cable was attached to a weight which was subsequently rotated around the axis of the connector to test it's mechanical robustness. This was of course also tested in other ways. One of the ways to test the durability of the product was to manually connect and disconnect the connector a lot of times. There was an automatic setup available, but it wasn't able to reproduce the exact effect from human handling. So it had to be done manually. That does sounds like a good summer job, doesn't it?

Some of the more awesome setups to watch were the radio rooms. These were too large rooms in the middle of the lab. The inside of one of the rooms looked like something you would get if you would cross a shark with a children's playpen. The whole room was covered in large foam spikes in order to dampen electromagnetic waves. This was used to test electromagnetic radiation from connectors, cables and antennas. One of the radio rooms took up a very prominent place in one of the electrical labs. There were large racks of measurement and analysis equipment around the chamber, to measure the electromagnetic radiation from the setup inside. In my opinion this chamber could also be used to play a very interesting game of marbles, or maybe a strange version of twister.

There were more tests that I could go into detail on, but I think this gives a good impression of the things I saw. The next and last part of my tour involved the workshop of the facility. This workshop is used to fabricate connectors and parts when the amounts required were less than they wanted to produce with mass production. My surprise however, was pretty big when I saw the workshop there. During my visit, one of the larger machines was busy extruding casings for connectors, when I asked how many connectors there were, I was told there would be about five thousand. When I think of manual production, I think of about three units. I guess I was wrong. There were machines in the workshop that could create metal parts, extrude casings and assemble these. I saw a few large molds for plastic casings, large sheets of metal being punched into the desired shape, and even a machine which assembled the parts of the product to complete system. I didn't even get to see the whole workshop, so I expect there to be a lot more. The next time I think of manual production, I guess I will see a large extrusion mold, punching out a few units per second, or something like that.

This marked the end of my tour. I had seen some of the internal workings of TE connectivity. However I expect that there is a lot that I haven't seen.

"When I think of manual production, I think of about three units. I guess I was wrong."

We quickly returned to the point of our meeting, where we had some time to look at the special system worked into the wall next to the door. There was a tablet like system that showed whenever the conference room was booked in an agenda. One could see how long the next meeting in the room would take and who would be there. A small thing to notice at a time like this, but is was told that they were busy with rebuilding parts of there facilities and this was one of the new innovations employed. I thought is was rather useful.

With that I had already reached the end of my tour at TE connectivity. I have seen some things that showed me how they work and operate and I hope that I have transferred some of my impressions of this interesting visit. I have learned something from it, and maybe you have too. There is of course a lot more going on at TE connectivity then the brief overview I have just given you, but I think that you will be able to find out more for yourself, if you want to.

$(AC/DC)^{-1}$

Author: Vera Nauta , Pascale van de Ven
Photos: Marcel Wenting

Solar energy has some advantages over energy from fossil fuels; solar energy does not cause any pollution in the form of carbon dioxide and unlike the supply of fossil fuels, the supply of solar energy will not reduce any time soon. This makes the project for the first year's students at the end of the second quartile very interesting. For the duration of the project they were provided with solar panels. The output of a solar panel is a DC voltage, which had to be fed back into the power grid, through a transformer. The input of the transformer had to be an AC voltage of 12 V effective.

To turn the DC signal into an AC signal, Pulse Width Modulation was used. Because we only used a positive voltage, it was only possible to build a rectified sine wave.

To turn this rectified sine wave into a full sine wave, an H-bridge was used. To make sure the phase of our signal corresponded to the phase of the grid, we used a comparator to check the polarity of the grid. The output of the comparator was sent to the microcontroller, which can adjust the PWM, and therefore the phase of the created sine wave.

> "This makes the project very interesting."

The solar panel will deliver the maximum power at a certain ratio between the voltage over the solar panel and the current through the solar panel. To find this maximum power point, the voltage and current are measured and sent to the microcontroller, which will calculate the power. The amplitude of the created sine wave will be varied and the effect on the power will be calculated. This way, the maximum power point can be found. The voltage over the solar panel is not high enough to create the desired amplitude. Therefore, a boost converter was added.

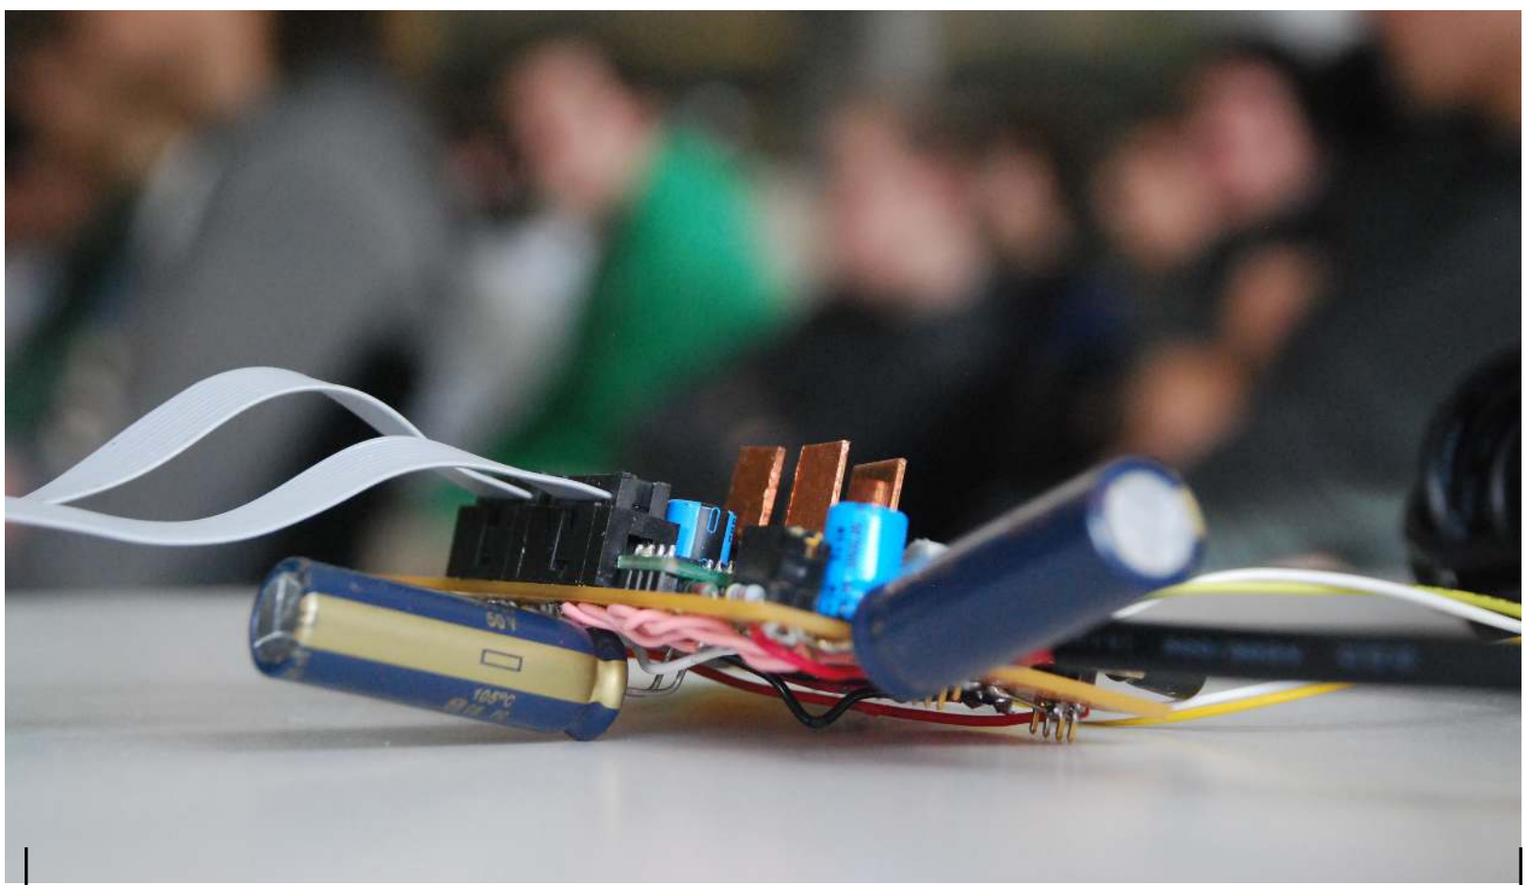

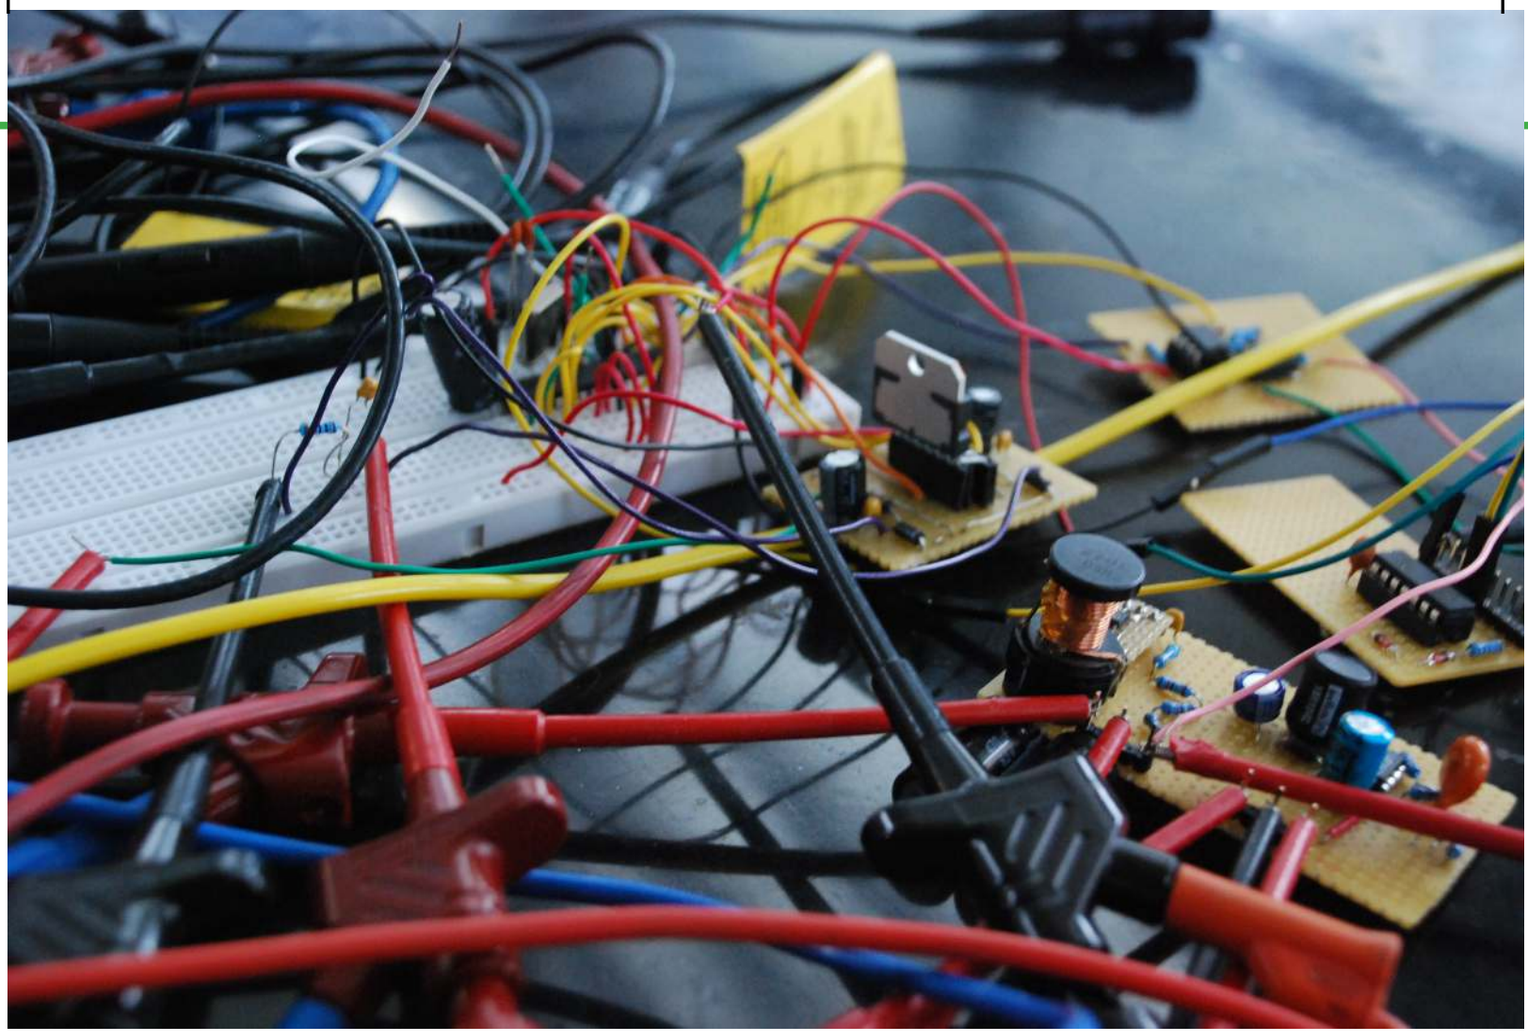

At the end of the project, all systems were tested. Even though our system looked like a total mess, and the maximum power point tracking did not work at all, we were one of the two groups who could actually deliver power to the grid. An overwhelming 1 W to be precise!

> "What we really liked about this project was that we were allowed to figure things out ourselves."

What we really liked about this project was that we were allowed to figure things out ourselves. Learning by trying to put theory into practice was a nice addition to the lectures. Some people complained that the project was too hard and that we had no idea what we were doing. The latter is certainly true, but it allowed us to learn a lot of new things. Of course it was also pretty awesome that we were actually able to feed power into the grid. Altogether, we had thee very entertaining weeks!

Chain Reaction

3D electronics 2B

Author: Anne-Johan Annema

Each edition of the Vonk has a main article written by one of the research groups at our University. Such an article discusses a recent development or trend in their field of research. It is interesting to listen to an outsider's comment on such a trend, being either positive or negative. For this purpose, starting this Vonk, we intruduce a new feature. As this text will be a critical glance on the previous main article, which in his turn is a comment on a technology or trend, we speak of a Chain Reaction. Anne-Johan Annema sets the ball rolling by commenting on TST's text on self-assembly.

ICO Integrated Circuit Design

It is frequently assumed that nowadays electronics are made in planar, two dimensional, IC technologies. It is true that CMOS (complementary Metal-Oxide-Semiconductor) transistors that are used in nearly every IC appear to be quite planar. However, already this very basic transistor is a true 3D structure. A cross section of a modern CMOS transistor is shown in figure 1 (from [1]). In the newest technologies that we can access nowadays at ICD the channel lengths are down to 28nm and decreasing to 22nm in about one year. The resolution on 28nm is about 1nm which translates into just a few atoms. The resolution in the vertical direction is even more scary: layers that are in the order of 1nm with a very low tolerance (gate oxide layers) and bulk layers with a thickness of about 10nm are available. Adding a relatively thick gate electrode turns this device into a small 3D-structure. Still we regard this MOS transistor as a planar device. Also finfets - which are basically 90 degrees out-of-plane-rotated MOS transistors - can be seen as quasi-planar devices while for marketing reasons they are frequently denoted as 3D-transistors.

> "The resolution on 28nm is about 1nm which translates into just a few atoms."

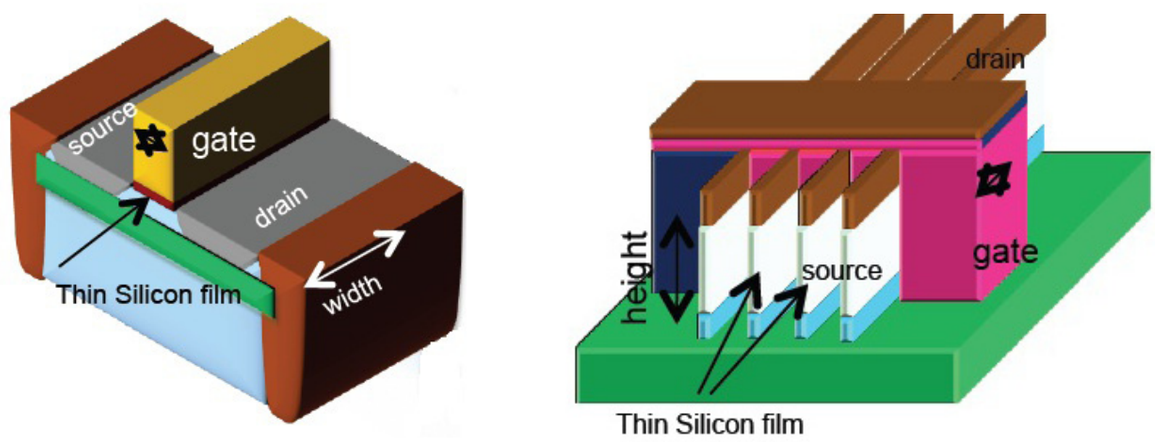

Figure 1: Modern MOS transistors [1]: almost planar ultra-thin-body transistors or just-planar-from-a-distance FinFETs

Nowadays 3D full custom electronics

The true 3D aspect in every modern and in every upcoming CMOS technology is in the interconnects between all transistors and passives in the electronic chip. This interconnection is required to get some non-trivial functionality from an electronic circuit, digital, analog, or mixed-signal.

Usually the interconnect is structured quite hierarchical. At a local (sub-circuit) level transistors and passives are interconnected via horizontal metal lines and vertical metal via's that are specifically designed to make the exact right connections between the various transistors that are usually individually sized, shaped and biased for optimum performance. The interconnect layers for local interconnections are located relatively close to the transistors, and are usually relatively short and narrow. This type of interconnect is equivalent to local roads inside the certain neighborhood in a town.

One level up in the hierarchy, the sub blocks with their many local interconnects and components are interconnected to power lines and to signal lines to and from other sub-circuits to form a circuit of subsystems. The interconnects for this purpose are usually thicker metal interconnect layers further away from the actual transistors. Just as any other layer and any component in a high performance electronic chip, this layer is specifically made to get exactly the intended functionality. The interconnects between the sub-circuits are similar to main roads in a town.

In modern systems, this procedure repeats itself one or two times to construct large electronic systems with huge computational power; analog, digital or mixed. The hierarchical approach is essential to handle the overall complexity of the system. At the same time very good interaction and cooperation between designers in all stages is required to get an overall good result. The semiconductor physics guys deliver the components - mainly transistors - that are sized, biased, used, pampered and squeezed out by the circuit designers to push performance. For analog and mixed-signal circuit designers, usually the goal is to approach or even circumvent the physical and practical limits to system performance of each technology. This clearly requires also a lot of knowledge of system level issues, or requires a lot of interaction with system level designers.

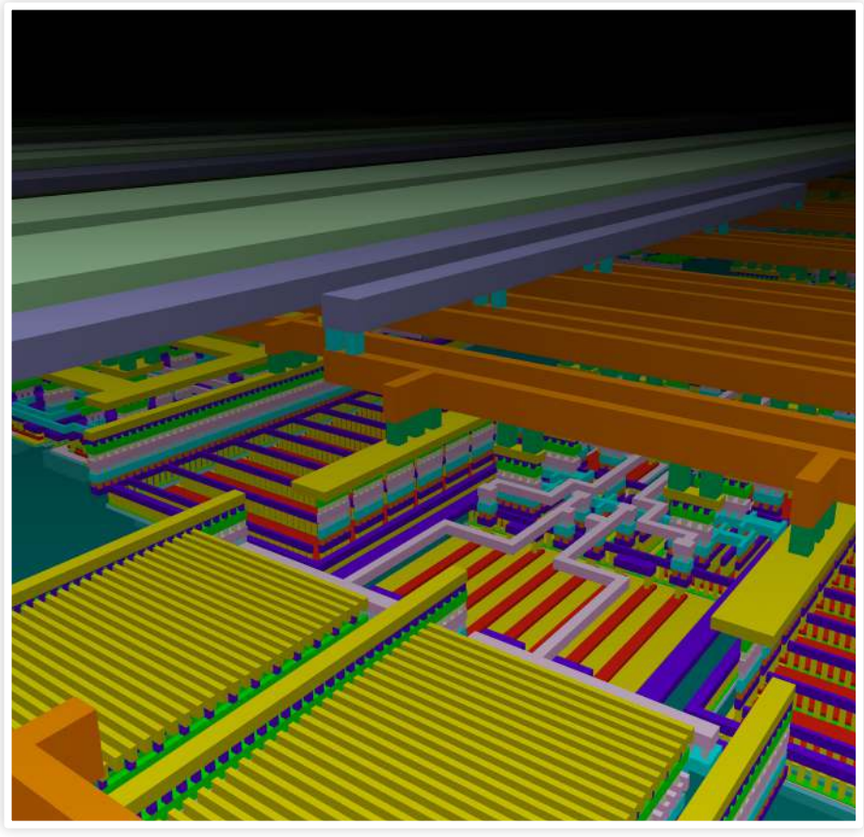

Figure 2: Left: 3D view of part of a chip design in 65nm CMOS technology.

> "For analog, digital and mixed signal electronics, self-assembled electronics are quite hard if not impossible."

As an illustration for the 3D construction of nowadays electronic chips, figure 2 shows a 3D view of a part of a chip design in 65nm CMOS technology made at the ICD group some years ago [2],[3]. In this view, the transistors are somewhere far down below and are hardly visible. The thin and short local interconnect lines can be seen in the lower layers, while the interconnect lengths and thicknesses increase towards higher layers. In the shown design, 8 metal interconnect layers are used, building a true 3D structure.

Figure 3 shows a cross section of a design in a newer 28nm process [4]. The transistors are located just above the buried oxide (BOX) layer and are not visible on this scale. Mainly visible are the interconnect layers that are thin and narrow at the bottom, for local interconnects, and are wider and more thick towards the top of the 3D structure for long distance signaling and power routing.

3D self assembling electronics

For analog, digital and mixed signal electronics, self-assembled electronics are quite hard if not impossible. This is because the electronics are fully customly designed,

both the transistors and the interconnects. Steered assembly just might be done using a Tetris-kind of approach which then would be quite time consuming, assuming that the very small transistors - nowadays down to about 5nm*30nm*60nm including drain and source region to contact - can be steered to the exact right position. Also the interconnects must be steered.

But don't get me wrong, the proposal in the previous Vonk is interesting: self-assembly might be used for very regular structures, with very regular active or passive devices and very regular interconnect. The system that satisfies this is a memory. Current advances in memories using transistors include both increasing the lateral packing density and decreasing the physical capacitor size for DRAM. The packing density is effectively tackled by Moore's law, which does not come cheap and which requires many research effort. The physical capacitor size in DRAM can be tackled in electronics and in e.g. making vertical capacitors: digging capacitors into the silicon.

3D self-assembly appears to be a disruptive technology that might push the memory density significantly. Still, the technology must include a non-self-assembled electronics part - fully customly designed - for read and write actions, including proper (regular) interconnects to these electronics. This can be built for example below the memory itself. For a cube memory matrix only custom designed read/write circuits are needed which is quite a bit less than for a conventional square memory matrix, assuming large memories.

Future of nowadays 3D full custom electronics

As explained previously, nowadays electronics are and will remain to be designed fully customly, to be able to implement high performance and well defined and non-trivial functionality. To be able to do this, all of nowadays chips are true 3D structures with (quasi) planar active devices (the transistors) and a huge stack of custom designed interconnections to route all supplies and signals. A next step towards even more 3D integration would be stacking a number of today's 3D structures on top of each other using extra vertical (through chip) via's for global power and signal routing. This would be feasible as the lateral dimensions of a chip are in the order of a few mm, while the actual height of the 3D stack is roughly 10um (containing tens of different layers, including the transistors and the many interconnect layers). Stacking tens or more of these structures could become feasible in the future which adds another dimension to our current 3D electronic chips.

Another nice thing about conventional 3D electronics is that the shrinking feature sizes of transistors will stop some day. Upcoming transistors (in a few years) will have gate lengths and interconnect feature sizes of about 10nm which is just about 20 atoms wide. Even then, transistor sizes are expected to continue to shrink for a few years. However, then the major increase in computational power and analog functionality must come from 3D integration of electronic systems, and from smarter circuit designs and smarter system designs. The end of the ever shrinking of transistors just may be the beginning of ever increasing challenges for circuit and system designers.

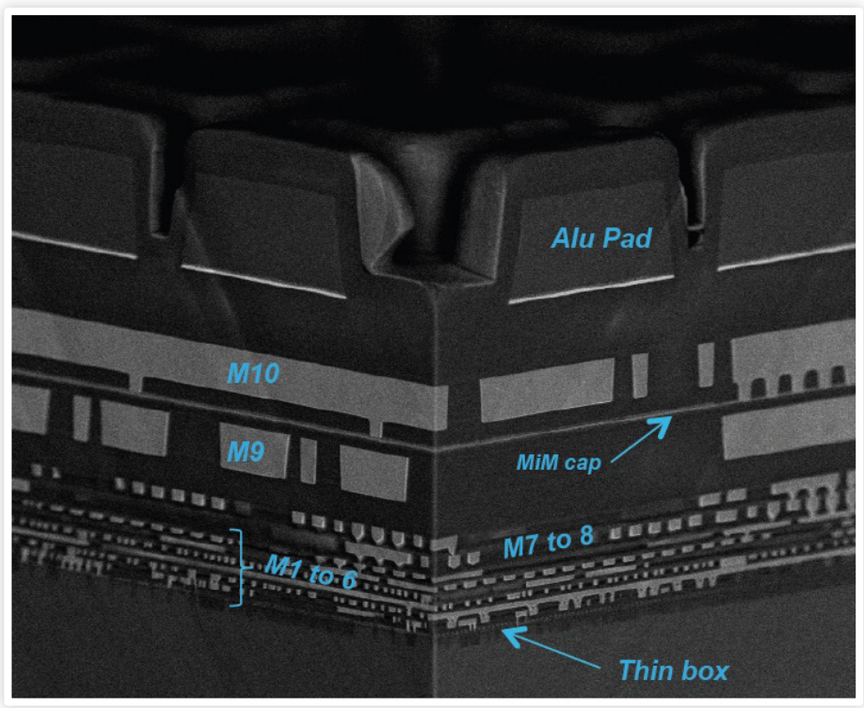

Figure 3: Cross section of a modern 28nm IC technology.

[1] J. Hartmann, "Planar FD-SOI Technology at 28nm and below for extremely power-efficient SoCs", Symposium: Fully Depleted Transistors Technology - December 10th, 2012 – San Francisco, CA

[2] M.C.M. Soer, screendump of part of his chip design

[3] http://sourceforge.net/projects/gds3d/

[4] L.A. Wolderink, P.J.Kamp and L. Abelmann, "3D self-assembled electronics", Vonk 1, pp 26

Cognitive radio

Author: Mark Oude Alink

Raar idee dat ik vier jaar geleden nog bezig was met afstuderen, en dat nu alweer mijn promotietraject erop zit. Mijn gecombineerde afstudeeropdracht ging over het verbeteren van een geïntegreerde spectrum analyzer, een apparaat waarmee de verschillende frequenties van signalen onderscheiden kunnen worden. Deze afstudeeropdracht voerde ik uit bij de informatica-vakgroep Computer Architectures for Embedded Systems (CAES) van professor Gerard Smit en bij de elektrotechniek-vakgroep Integrated Circuit Design (ICD) van professor Bram Nauta.

Al halverwege mijn afstuderen kwam naar voren dat er een project aan zat te komen waarin de belangrijkste doelstelling was: 'het in kaart brengen en oplossen van problemen om een Cognitive Radio (CR) in goedkope en low-power hardware te realiseren'. Dit onderzoek paste behoorlijk goed bij mijn afstudeeronderwerp en zou bij dezelfde vakgroepen plaatsvinden.

Ik wist niet meteen of ik dit wel wilde; is het niet beter je blik te verruimen, na ruim zeven jaar aan deze universiteit te hebben gestudeerd? Voor alles zijn voor- en nadelen te bedenken, maar tijdens mijn afstuderen was ik er al wel achter gekomen dat onderzoek doen mij wel ligt. De keuze om AIO te worden was dus gemaakt. Gezien mijn brede interesse paste de CR opdracht goed bij mij, omdat het veel vakgebied-overschrijdende onderwerpen kent. Een voordeel was ook dat ik al een soort vliegende start had door mijn afstuderen, wat zeker gunstig is gezien de ervaring van veel AIO's: het eerste jaar ben je vooral aan het inlezen en na vier jaar ben je veel minder ver dan je zou willen. Ondanks de vliegende start ben ik toch minder ver gekomen dan ik van tevoren had gedacht. Dat neemt echter niet weg dat het enorm goed bevallen is, want je hebt een vrijheid die je in het bedrijfsleven niet snel zult tegenkomen. Zowel qua onderwerpen waar je aan werkt als hoeveel tijd je daaraan besteedt, zonder al teveel strikte deadlines.

> "Na vier jaar ben je veel minder ver dan je zou willen"

Cognitive Radio

Om te begrijpen wat CR inhoudt, moeten we eerst weten dat draadloze toepassingen op dit moment elk hun eigen stukje spectrum krijgen toegewezen (zie figuur 1), eventueel na betaling, maar zeker pas na uitgebreide tests. Ook zijn er enkele zogeheten Industrial Scientific and Medial (ISM)-banden, waar elk protocol na goedkeuring gratis gebruik van mag maken. De meest bekende is de ISM-band op 2.4GHz, waar onder andere WLAN, Bluetooth en ZigBee in opereren. Deze vrij beschikbare bandbreedte is echter zeer beperkt, wat zich praktisch uit in een gelimiteerde snelheid en een gelimiteerd aantal verschillende kanalen in een beperkte ruimte voor WLAN.

Door deze toewijzing is spectrum zo schaars geworden, dat bedrijven grof geld betalen voor een paar MHz bandbreedte. Om een voorbeeld te noemen, in 2008 heeft de

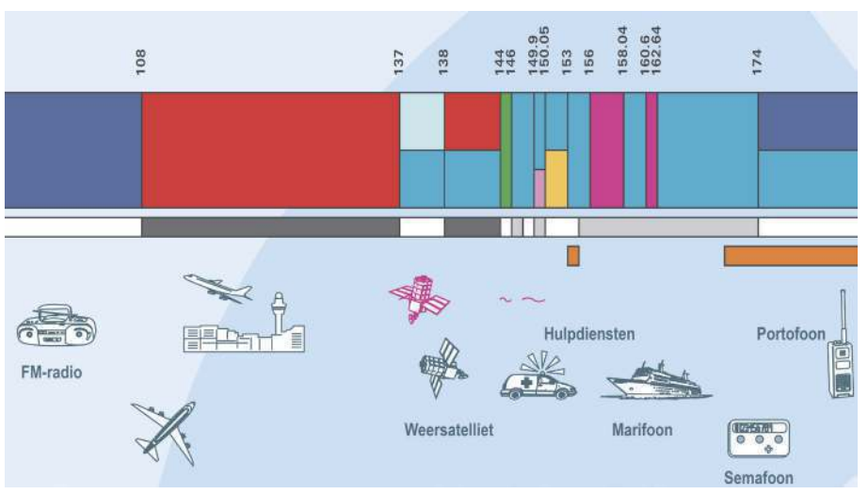

Figuur 1: Versimpelde frequentie-indeling tussen 100MHz en 200MHz in Nederland
Bron: Agentschap Telecom

FCC (de toezichthouder op het spectrum in de VS) een veiling gehouden waarin een jaar exclusief gebruik van 52MHz bandbreedte geveild is voor 19.6 miljard dollar. Voor het spectrum in de buurt van de 700MHz werd dus 400 miljoen dollar per MHz neergelegd! Metingen hebben echter uitgewezen dat het gebruik van de statisch toegewezen frequenties sterk varieert. Gemiddeld genomen is het spectrum op een willekeurige plaats en tijdstip slechts voor 3% tot 20% in gebruik. Hier valt dus nog een hoop winst te halen!

De term "Cognitive Radio" is in 1998 bedacht door Joseph Mitola III, die in 2000 op dit onderwerp promoveerde, maar heeft in de loop van de tijd verschillende betekenissen gekregen. IEEE definieert het als volgt: "A type of radio in which communication systems are aware of their environment and internal state and can make decisions about their radio operating behavior based on that information and predefined objectives." Cognitief slaat dus op een mate van intelligentie en adaptiviteit. In de meest ruime zin van het woord bestrijkt CR zich over zeer veel wetenschappelijke gebieden, maar in ons project beperken we 'environment' tot het frequentiespectrum en 'operating behavior' tot het uitzenden en ontvangen op beschikbare frequenties.

Enkele onderzoeksrichtingen binnen CR

Zelfs binnen dit beperkte kader, waarin de CR het spectrum scant op zoek naar vrije ruimte om in te zenden/ontvangen, zijn er vele niet-triviale problemen. Ik noem er hier een paar, zonder volledig te willen zijn.

Een vrij zwak TV-signaal, dat nog net sterk genoeg is om ontvangen en gedecodeerd te worden door een TV-ontvanger, komt door een muur (dus nog een keer sterk verzwakt) bij een CR aan. Betekent dit nu dat dit stukje spectrum vrij is of niet? Want hoe weet een CR of er een TV-ontvanger in de buurt is? En hoe weet de CR of het signaal door een muur heen ontvangen wordt of dat het een directe line-of-sight is? Dit is een soort hidden-node problem, maar dan in een nieuwe context.

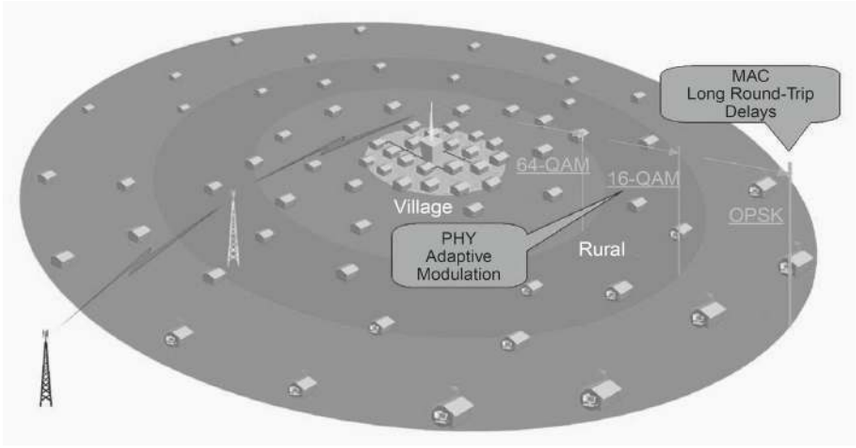

Figuur 2: IEEE 802.22 heeft als doel om landelijke gebieden van breedband-internet te voorzien zonder al te dure infrastructuur door CR-technologie in de TV-banden te gebruiken. Bron: Cognitive radio communications and networks - principles and practice

Hoe wordt de beschikbare bandbreedte verdeeld tussen verschillende CRs? Onderling communiceren, via een master-slave systeem, of gewoon wie het eerste komt het eerste maalt? En hoe beginnen twee CRs überhaupt met communiceren? Het spectrum bij de ene CR kan significant anders zijn dan bij de andere CR. Moet er een gemeenschappelijk control channel zijn? Maar dat betekent weer statische toewijzing. Is dat wel wenselijk, en zo ja, hoe regel je dat op wereldniveau?

Het scannen van het spectrum en het beslissen welke band beschikbaar is, danwel het beste is qua eigenschappen, valt niet binnen één laag van het OSI-layer model. Hoe moet daarmee omgegaan worden? Sommige gelicenseerde gebruikers hebben een (redelijk) vast patroon van frequentiegebruik. Wat voor vorm van kunstmatige intelligentie of patroonherkenning is nodig om dat detecteren, te voorspellen en uit te buiten?

Onderzoek naar hardware voor CR

In ons project wordt er vooral gekeken naar het ontwikkelen van hardware om zo'n CR mogelijk te maken; we maken ons niet druk om veel van de eerder genoemde problemen, hoewel ze natuurlijk erg relevant zijn. Het maken van de hardware heeft echter weer zijn eigen problemen, waarvan ik hier de belangrijkste noem.

Bij statische frequentie-toewijzing worden speciale hogekwaliteitsfilters gebruikt om te zorgen dat de zender nauwelijks vermogen buiten zijn frequentiebandje uitzendt, omdat dat anders andere gebruikers zou storen. Deze filters zijn echter duur en totaal niet flexibel. Dit is ook een reden dat bijvoorbeeld een quad-band telefoon een stuk duurder is dan een dual-band telefoon. Hoe moet een flexibele CR dit oplossen? Het is relatief duur en onpraktisch om voor elke mogelijke band een apart filter te hebben. Er zijn wel allerlei digitale filter-technieken mogelijk, maar dat kan veel rekenkracht kosten en daarnaast moet het uiteindelijk alsnog door een niet-ideaal analoog circuit.

Hoe maak je een spectrum analyzer on-chip die niet €50.000 euro kost en 200W verbruikt, zoals een moderne spectrum analyzer in het lab, maar die nog wel de benodigde performance levert? Waar kun je op beknibbelen en waar niet? Wat voor processing voer je uit om een beslissing te nemen of een frequentieband vrij is of niet? Hoe zorg je ervoor dat je dat met een energiezuinige processor, die je waarschijnlijk toch al hebt voor andere toepassingen van je (mobiele) apparaat, kunt behappen?
Hoe realiseer je bovenstaande functies in een modern CMOS-proces waarin je slechts ongeveer 1V ter beschikking hebt en waar de transistoren voor digitaal weliswaar erg goed zijn, maar voor analoog niet echt? En hoe zorg je dat de totale powerconsumptie laag genoeg blijft, zodat je niet elke tien minuten je accu hoeft op te laden?

Techneuten & Regelgevers

Een erg interessante discussie is die tussen techneuten en regelgevers. De regelgevers willen graag de gelicenseerde gebruikers van het spectrum beschermen, aangezien die vaak grof geld hebben betaald om hun diensten te kunnen aanbieden. Om garanties te kunnen geven dat hun diensten niet gestoord worden, moet een CR zelfs bij hele zwakke signalen deze band al als "in gebruik" markeren.

In verband met allerlei mogelijk effecten die optreden als radiogolven reflecteren en interfereren (bijvoorbeeld fading en shadowing), kan de CR een veel zwakker signaal zien dan een naburige ontvanger van de diensten. Dit betekent dat zwakke signalen gedetecteerd moeten worden, ver onder het ruisniveau van elke ontvangers. Hoe doe je dat als je niet weet wat voor soort signaal het is?

Techneuten vinden daarom veelal dat de limieten omhoog moeten. Televisiezenders zijn daar weer niet blij mee, omdat er dan kans bestaat dat de signaalkwaliteit bij hun abonnees omlaag gaat door de storing veroorzaakt door CRs. Dit soort discussies heeft in 2004 zelfs geleid tot het oprichten van de DySPAN-conferentie, een conferentie waarin regelgevers, onderzoekers en praktiserende bedrijven bij elkaar komen om face-to-face over dit soort zaken te kunnen redetwisten.

Voordelen van CR

Afgezien van efficiënter gebruik van het spectrum, hebben wij als consumenten iets aan CR? Jazeker! Niemand vindt het leuk dat het rond de jaarwisseling niet of nauwelijks lukt om een SMS'je te versturen, hoewel daar natuurlijk nog wel mee valt te leven. Met standaard GSM zijn er simpelweg te weinig kanalen beschikbaar om iedereen te bedienen, terwijl met CR-technologie veel meer kanalen gevonden kunnen worden.

Zeer belangrijk daarentegen is de mogelijkheid om hulpdiensten betrouwbaarder en breedbandiger communicatie te bieden. CR maakt het mogelijk dat de vele medewerkers niet meer gelimiteerd zijn door de 10MHz bandbreedte dat C2000, het huidige systeem, biedt. Dan waren er bij gebeurtenissen, zoals bij de Vuurwerkramp in Enschede in 2000 of de aanslag op Koninginnedag in 2009, geen problemen geweest door overbelasting. Ook staat C2000 bekend om zijn verbindingsproblemen in gebouwen, terwijl een CR de mogelijkheid biedt op een andere frequentie over te schakelen waar wel bereik is. De grotere bandbreedte maakt het ook mogelijk om foto's en/of videobeelden te kunnen overzenden, wat een grote toegevoegde waarde kan hebben.

Een ander voordeel is het mogelijk maken van breedband internettoegang in landelijke gebieden. Voor Nederland is dat niet zozeer een probleem, maar in de VS heeft slechts 51% van de bevolking breedband internettoegang, omdat 75% van de bevolking op 2% van het landoppervlak leeft. Het is veel te duur om de infrastructuur aan te leggen om iedereen van breedband internet te voorzien. Echter, de frequentieband van 50-860MHz, waar de TV-kanalen zitten, biedt zeer goede propagatie-eigenschappen, waardoor het eenvoudig is om draadloos meer dan 30km te overbruggen. Met CR wordt het dus mogelijk om met weinig investering grote gebieden van breedbandinternet te voorzien, zie figuur 2. Dat is de reden dat zelfs bedrijven als Google en Microsoft op de DySPAN-conferentie aanwezig zijn om hun eigen CR-hardware onder de aandacht te brengen.

Recentelijk is de IEEE 802.22 standaard in het leven geroepen voor precies deze toepassing, maar dat is niet de enige ontwikkeling. Ook de 802.11 (WiFi) familie krijgt een nieuwe telg, 802.11af, die in de TV-banden opereert. Je zou hiermee bijvoorbeeld simpelweg een snellere WLAN-verbinding in je huis of op de campus kunnen krijgen. British Telecom heeft zelfs al ideeën om met deze technologie de straten van hele wijken in Londen van draadloos internet te voorzien, zie figuur 3.

Conclusie

Samenvattend is CR dus een nieuw paradigma voor draadloze communicatie met vele mogelijkheden, waarvoor echter eerst nog vele problemen overwonnen moeten worden alvorens het een succes kan worden. Ons project heeft een aantal van deze problemen verminderd of opgelost. Wellicht dat we deze in een toekomstige Vonk kunnen laten zien.

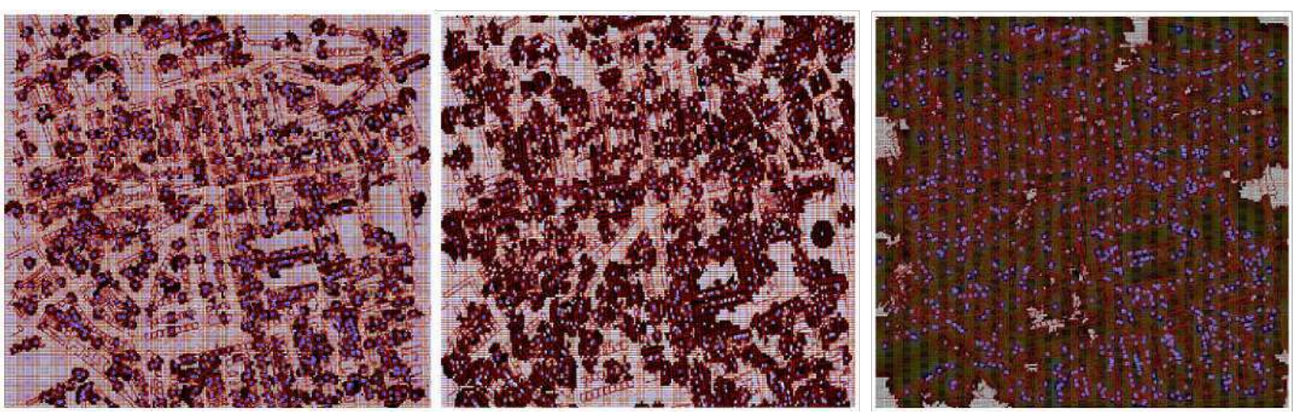

Figuur 3: WiFi-dekking (simulaties) in Notting Hill, Londen op 5GHz, 2.4GHz en in de TV-band.
Bron: British Telecom

Activatelunch

Kerstdiner


Scinterklaas
MILFBA

Finger vein recognition

Signals and Systems Group

Author: Fieke Hillerström

Signals & Systems

Biometrics are becoming increasingly important for recognition of people. Biometrics refer to the recognition of individuals by their characteristics or traits. Fingerprints and face recognition are traditional methods. Among the recently new recognition methods are the ones focused on finger vein patterns. For the past quartile I have worked on my Bachelor assignment at the Signals and Systems group, focused on finger vein recognition. This assignment proceeds the Master assignment of Bram Ton, who constructed a database of finger vein images. My assignment was to determine the performance of local binary patterns on finger vein recognition. Next to that I looked at the performance of the finger shape as biometric classifier.

Finger vein recognition is a relative new method in the field of biometrics. It uses the crisscrossing vein patterns in a human finger to identify individuals. An example of a finger vein image is shown in figure 1. One of the advantages with respect to the more traditional biometrics technologies is the difficulty of counterfeiting the finger vein patterns. Because the veins are hidden inside the body, the surface conditions of the hands have no influence on the authentication process.

To capture the finger vein images, infrared light is needed, because the patterns are located under the skin. Infrared (IR) light can penetrate relatively deep into the skin and passes through human body tissues, but is absorbed by the hemoglobin in blood. To capture the finger vein images, IR illuminators are placed above the finger such that the light penetrates the finger. An infrared camera is used to capture the images.

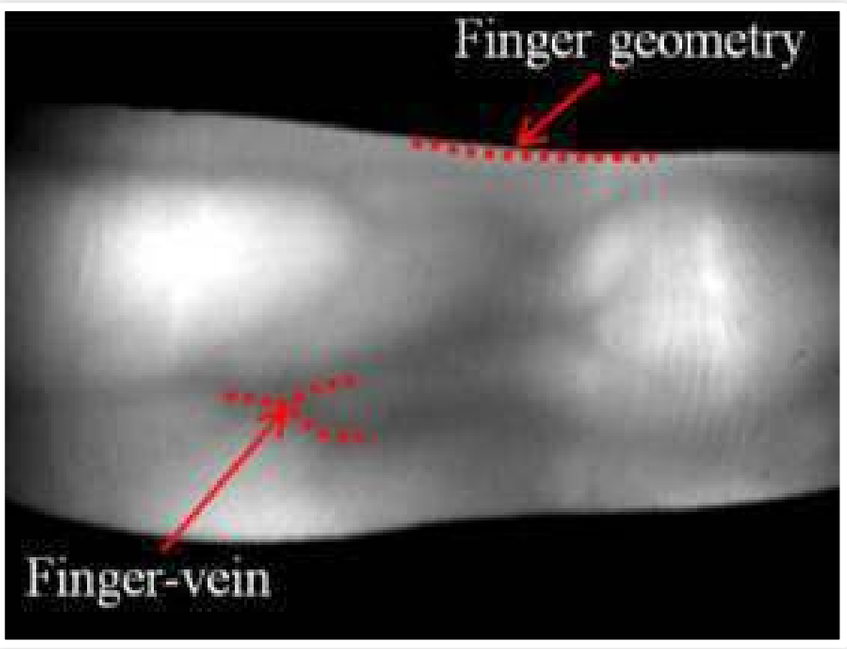

Figure 1: Example of a finger vein image

An example of a device to capture them is shown in figure 2.

My assignment was to implement local binary patterns (LPBs) on the finger vein images. LPBs are used as a method for texture analysis. Basically it is a binary code describing the local texture pattern, built by thresholding a neighborhood by the gray value of its center. Every pixel in a picture is compared to the values of its neighborhood, which results in a label for this pixel. The outcome of this thresholding results in a binary code, which is converted to a decimal value, as shown in figure 3. The advantage of LPBs are their simplicity and their independency with respect to light exposure.

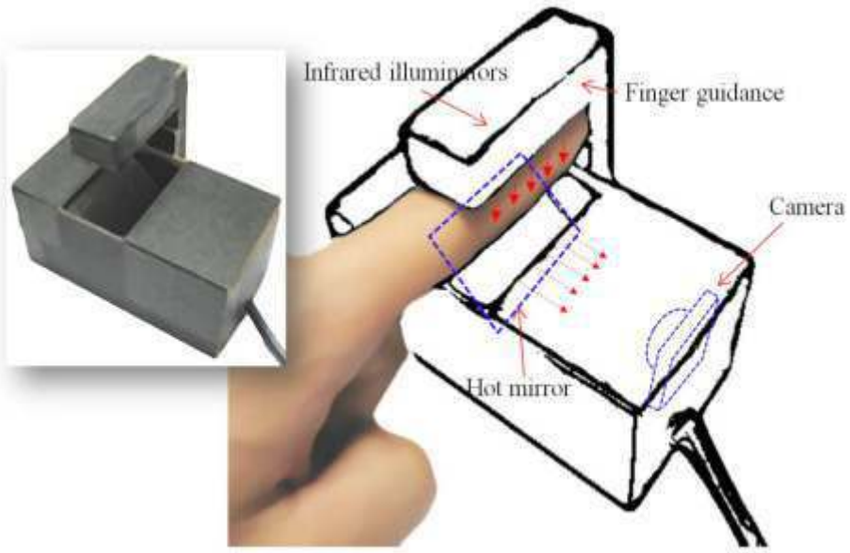

Figure 2: Capturing finger vein images

"Using the finger shape descriptors leads to better performance"

All these LBP values are combined in a histogram, which is used as a texture descriptor. This histogram contains information about the distribution of patterns over the whole image, but loses spatial information. To retain the spatial information, the image could be divided into sub regions over which a new histogram is calculated. This histogram is a measure for the amount of different structures in an image. One of the advantages of LBPs is that the method is less dependent on different light exposure in two sample intervals.

Before calculating the LBP histogram for a finger vein image, every image has to be normalized. The finger images need to have the same size and orientation. After this normalization, the local binary pattern histograms are calculated and the edges of the finger image are used as finger shape descriptor. These two classifiers are combined and the performance is determined.

From this research, it can be concluded that using the finger shape descriptors leads to better performance than using the local binary patterns. But the local binary patterns could be promising in combination with more complicated recognition algorithms.

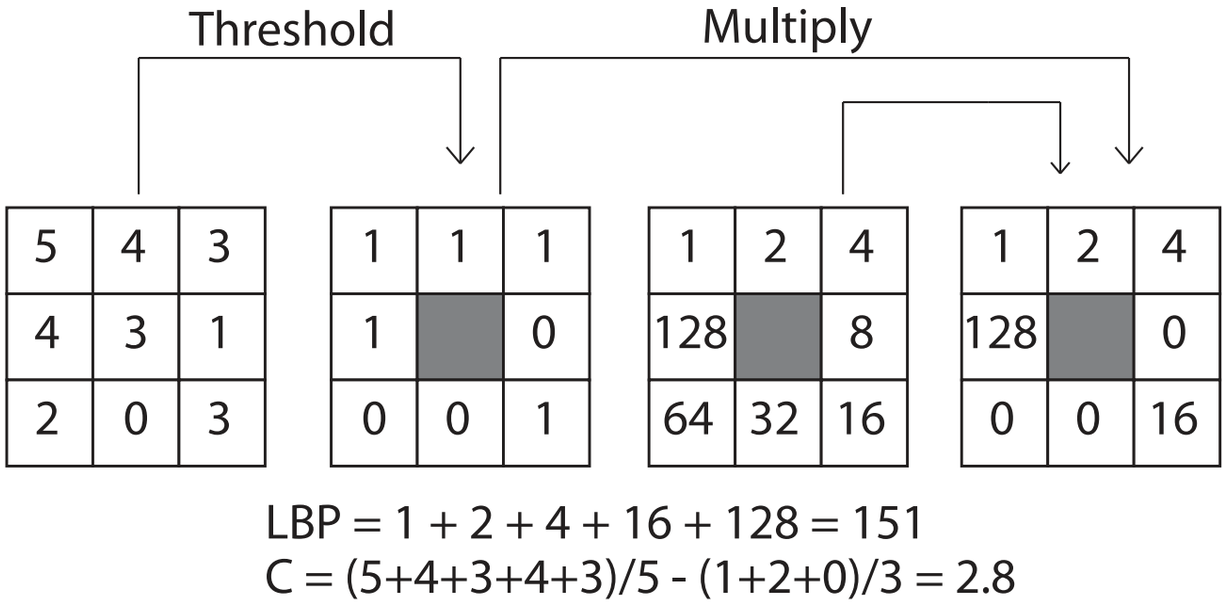

Figure 3: The process of calculating the local binary pattern of a pixel

The Logical Volume Manager

Partitions made easy

Author: Koen Zandberg

Hard disk partitions can be a pain to manage correctly. Sometimes you have a partition that is too small and it needs to be enlarged, or it is too large and you're wasting precious space. Unfortunately partitions aren't flexible and can't easily be resized, even if it is possible it takes hours to do. These problems can be circumvented by using the Logical Volume Manager, if you have Linux.

What is the Logical Volume Manager?

The Linux Logical Volume Manager is disk management software. It acts as a sort of vitalisation layer between the partitions and the physical hard disks. The Logical Volume Manager can manage flexible partitions which can be easily configured and changed. The Logical Volume Manager can create, resize and destroy partitions with a few simple commands. For example, one can, when installing the operating system create a few partitions, but only as large as necessary when installing. When it needs more space on this partition, it can extend the partition and continue working. Normally, a partition needs to be one continuous space, but the LVM circumvents this problem. This can all be done on a live system and doesn't require a reboot.

How does it work?

The Linux Logical Volume Manager works by utilizing the device mapper kernel module. The device mapper is a generic framework for mapping one block device (for example a hard disk) to another (virtual) block device. It forms the foundation of many hard disk technologies in Linux. The LVM makes use of it, as do software (fake) RAID and encrypted volumes. The device mapper is used to map the virtual partitions created by the LVM to space on a real hard disk.

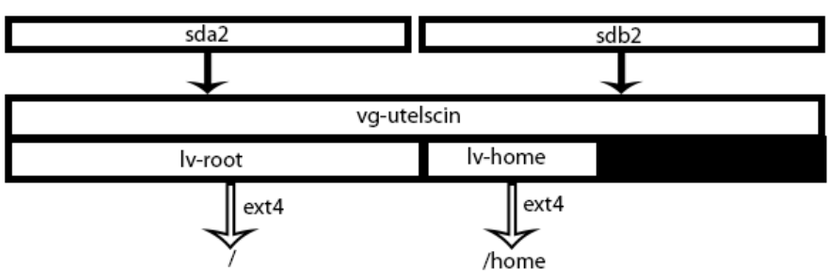

Figure 1

The LVM works as follows. The LVM needs a partition to work with, this is a so called physical volume or PV. This is the space the LVM can use to store data. Physical volumes form the link between the actual hard disk and the Logical Volume Manager.
Physical volumes are grouped together in a volume group or VG. The volume group is the main working unit of the LVM. It manages the physical volumes it consist of and with that the total space of the hard disk pool.
From the volume group we can create logical volumes (LV), which can be used as partitions for a system. The logical volumes can be accessed as real disks. A logical volume can be formatted to have a file system on it and is the resulting storage block of the Logical Volume Manager.

Each physical volume in a volume group has some space reserved to store the entire layout of the volume group. Thus if a physical volume goes missing, for example it breaks down, the entire layout of the volume group

```
root@utelscin:~$ fdisk -l

Disk /dev/sda: 250.1 GB, 250059350016 bytes
255 heads, 63 sectors/track, 30401 cylinders, total 488397168 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x000910f0

   Device Boot      Start         End      Blocks   Id  System
/dev/sda1   *        2048      194559       96256   83  Linux
/dev/sda2          194560   488396799   244101120   8e  Linux LVM

Disk /dev/sdb: 250.1 GB, 250059350016 bytes
255 heads, 63 sectors/track, 30401 cylinders, total 488397168 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x00096769

   Device Boot      Start         End      Blocks   Id  System
/dev/sdb1            2048      194559       96256   83  Linux
/dev/sdb2          194560   488396799   244101120   8e  Linux LVM
```

Figure 2

can still be recovered. The rest of the physical volume is divided into small parts between 8 KB and 512 MB called Physical Extends (PE). These form the basic building blocks of a logical volume and are used to store your data.
A Logical volume exists of Logical Extends (LE). These map directly to one or more physical extend on the physical volume. This mapping between logical extends and physical extends is what the Logical Volume Manager gives its flexibility. A Logical volume can appear as one continuous space to the operating system, but can in fact be spread across multiple disks. The data is no longer bound to a specific area of the hard disk.

With this mapping, we can build a logical volume from a volume group. We can map a number of logical extends to a number of physical extends and create a partition of a certain size. The physical extends don't have to be contiguous and can be from any disk of the volume group. This feature indirectly enables us to make partitions larger than a single disk of a volume group. If we need a larger volume, we can instruct the LVM to add more extends to a logical volume. If it can do with less space we could instruct the LVM to remove some empty extends from the volume and shrink it.
With these tools we can resize any volume to fit our size requirements. Afterwards we only have to notify the file system on it that the volume under it has changed. A few file systems support live resizing and can adapt to the changes of the underlying volume. There are no limitations in the number of logical extends a logical volume can have. When given enough disks, a volume group could grow indefinitely. There's no IO penalty on using larger or smaller extends, only the tools for managing the logical volumes could suffer from a penalty. We could use larger extends, but a logical volume can only grow in steps of the size of an extend. So if we pick 512 MB as an extend size we can only grow our volume in steps of 512 MB. For most systems the default size of 4 MB is suitable.

> "When given enough disks, a volume group could grow indefinitely."

Basic usage

If we want to use the Logical volume manager we first need a partition we can use as physical volumes. In the case of Utelscin, one of Scintilla's servers, we have two disks with each a separate partition for our physical volume (figure 2). We have sda2 and sdb2 as our physical volumes. Before we can use these as a physical volume we have to format them as physical volumes. This can be done with "pvcreate". For example: "pvcreate /dev/sda2" creates a physical volume on the partition sda2. As visible in figure 2, the partition also gets Linux LVM as a label so it can be identified as a physical volume. We can ask our systems which physical volumes it had with "pvdisplay". An example of this is given in figure 3. As we can see we have 2 physical volumes and we can view some statistics about these volumes.

Now that we have our physical volumes, we can create a volume group with our physical volumes. With "vgcreate vg-utelscin /dev/sda2 /dev/sdb2" we create the volume group named "vg-utelscin" consisting of the physical volumes sda2 and sdb2. The name of the volume group is used as an identifier for the group and is used in most commands to specify which group the command applies to. Therefore it makes sense to use either a prefix or a suffix to identify it as a volume group, in this case the "vg-". Also here we have a command to display statistics, not surprisingly called vgdisplay. The result of this for Utelscin can be seen in figure 4. We can see that the volume group has 2 active physical volumes and that there's about 150 GB of free space remaining. This is space that is not yet allocated to a logical volume. Do not confuse this with free space on partitions.

From our new volume group we can create logical volumes. The creation of the volumes is the most complicated command of the LVM. This is because there are many options to select from for more advanced configurations. If we want to create a simple volume for our root partition we can use the following command: "lvcreate -L5G -nlv-root vg-utelscin". This creates a new logical volume of 5GB named lv-root in the volume group vg-utelscin. As with the volume groups, we used a prefix to identify the logical volume as one. There's now a partition that can be used for the system. However, before it can be used it needs a file system. Thus it needs to be formatted with a file system, for example ext4 with "mkfs.ext4".
If our 5 GB is not enough we can extend the volume with another 5 GB to make it 10 GB in size with "lvextend -L+5G /dev/

```
root@utelscin:~$ pvdisplay
  --- Physical volume ---
  PV Name               /dev/sda2
  VG Name               vg-utelscin
  PV Size               232.79 GiB / not usable 4.00 MiB
  Allocatable           yes
  PE Size               4.00 MiB
  Total PE              59594
  Free PE               19081
  Allocated PE          40513
  PV UUID               jA6DnK-rXK6-V8Dp-qScH-PBZi-2uv3-MLRyez

  --- Physical volume ---
  PV Name               /dev/sdb2
  VG Name               vg-utelscin
  PV Size               232.79 GiB / not usable 4.00 MiB
  Allocatable           yes
  PE Size               4.00 MiB
  Total PE              59594
  Free PE               19087
  Allocated PE          40507
  PV UUID               VN99og-v5pn-fJXw-uI5N-cKZ2-5IO2-KXqjVf
```

Figure 3

vg-utelscin/lv-root". ext4 supports live resizing with resize2fs so we don't have to take down this partition to resize it. Of course as with physical volumes and logical volumes we can view statistics about our new logical volume with lvdisplay. In figure 5 we can see the result of "lvdisplay /dev/vg-utelscin/lv-root". On Utelscin this volume is currently 22 GB and is mapped to 2 physical volumes. There are many more commands for managing your volumes. All of these starting with pv for operations on physical volumes and vg and lv for volume groups and logical volumes.

> "Many configurations are possible."

More possibilities

With the possibilities of the device mapper, we can do much more than just map one logical extend to a random physical extend. Many configurations are possible.
Since the LVM can use any block device for its partitions it is possible to stack multiple layers on top of each other. A lot of setups use a RAID configuration with on top of that a physical volume for added flexibility. Other possibilities include using a encryption layer with a physical volume for flexibility and creating multiple partitions on top of the encryption layer.
But the LVM can do some nice tricks on its own too. Most of these trick use a creative mapping of the logical extends to the physical extends.
Lets assume we have a disk which is starting to fail. Obviously we want to swap it for a new healthy drive but we don't want any downtime or nasty partition copying. We can add the new disk live to our server. We can then create a new physical volume on this disk and let the LVM move all the physical extends of the old drive to our new drive. The LVM will first copy a physical extend to the new drive. Then it changes the mapping of the logical extend to the physical extend. This way the operating system doesn't notice any downtime or any changes to the logical volume.
Another option is that we can map the logical extends of a logical volume each to two physical extends on different physical volumes. This creates a logical volume that is mapped to two physical volumes. This way, the data is stored on two physical disks, this way if one of the disks fails we still have our data on the physical extends of the other disk. Readers with a bit of knowledge in this field will recognise this as a RAID1 setup or mirroring. The LVM is able to do a number of operations on these volumes. We can while staying live split of a mirror and create a new logical volume with it. We can even split it off, track everything that changes while being split of. Later we can merge the separate volumes together and synchronise everything that has changed in the meantime. This can be used to split off a complete partition, hand the hard disk to someone. He then can change a few files or add something and merge it back into your disk array, letting the Logical Volume Manager handle the syncing between the disks.
An other option is to alternatively map logical extends between disks, mapping logical extend 1 to physical extend 1 on physical volume 1, mapping LE2 to PE1 on PV2

```
root@utelscin:~$ vgdisplay
  --- Volume group ---
  VG Name               vg-utelscin
  System ID
  Format                lvm2
  Metadata Areas        2
  Metadata Sequence No  165
  VG Access             read/write
  VG Status             resizable
  MAX LV                0
  Cur LV                24
  Open LV               22
  Max PV                0
  Cur PV                2
  Act PV                2
  VG Size               465.58 GiB
  PE Size               4.00 MiB
  Total PE              119188
  Alloc PE / Size       81020 / 316.48 GiB
  Free  PE / Size       38168 / 149.09 GiB
  VG UUID               1XVrJ7-29fi-rNC2-JSnj-F0bf-soxh-a7fM0z
```

Figure 4

```
root@utelscin:~$ lvdisplay /dev/vg-utelscin/lv-root
  --- Logical volume ---
  LV Path                /dev/vg-utelscin/lv-root
  LV Name                lv-root
  VG Name                vg-utelscin
  LV UUID                O9CisX-eUxM-pXB1-KoXm-xOP8-acmR-n1eXCb
  LV Write Access        read/write
  LV Creation host, time ,
  LV Status              available
  # open                 1
  LV Size                22.00 GiB
  Current LE             5632
  Mirrored volumes       2
  Segments               1
  Allocation             inherit
  Read ahead sectors     auto
  - currently set to     256
  Block device           253:4
```

Figure 5

and so on. This way the data is spread over 2 disks so we can read from both disks when requesting files. This is close to a striped setup, also known as RAID0. Because the operating system can now read its data from two disks simultaneously, the data transfer rate is theoretically doubled.
The LVM can implement most RAID setups by mapping the logical extends in a certain way to the physical extends. More complex arrays like RAID5 and other configuration with parity are possible as well.

> "More complex arrays with parity are possible as well. "

The LVM has a copy on write system implemented. Copy on write is a data write policy. It doesn't change data when it is modified, but instead writes a new block of data to the disks and changes the pointer to it. The old data is kept on the disk until it is overwritten by some other data. This enables a few more features for the LVM. Two things that are also possible are snapshotting and volumes with an virtual size.
Snapshots are made by creating a new volume. The logical extends of the new volume are mapped to the same physical extends of the volume the snapshot is created from. When something changes to one of the two volumes, the copy on write functionality is used to make a new physical extend for the changed extend. The changed extend maps the logical extend to the new physical extend, but the other logical volume keeps the old physical extend. This way a snapshot is only as large as the changed logical extends. With this we could make a backup that is only as large as the data that has changed in the meantime, growing with every change. We could use this to make consistent backups of our disks. Create a snapshot when we want a backup and backup this snapshot. We can be sure that this snapshot doesn't change while the original copy can still be used in a live system.

Volumes with a virtual size are volumes that have a size that is larger than the actual size. We can create a volume that is 5 GB large, but at the time of creation only has 1 GB of actual allocated space. This is created by mapping the logical volume to two devices. One that is mapped to the actual disk and one that is mapped to a virtual device that returns zero's of data when read from, appearing empty to the file system on top of it. When we need more space we can expand our actual volume without having to extend the file system. The file system already assumes it is 5 GB.

The logical volume manager gives us a set of tools to manage our file systems in an efficient and flexible way. It enables us to look at file system management in a different way, no longer limited by the standard problems with file systems.

A Brief history of LVM

LVM is already quiet a few years old. The current implementation of LVM is still being developed. But the basic idea of LVM, granting system administrators and users some flexibility with partitions is almost 20 years old. Netware 4 (1993) was already able to do basic LVM-like operations like resizing partitions. Even the tape label standard (ANSI X3.27 1970-1987) has the ability to make one large partition out of multiple smaller physical tapes.

The basic idea of the LVM as we know it nowadays comes from Heinz Mauelshagen. In 1998 he designed the first version of LVM, now know as LVM1, based on the LVM found on the HP-UX operating system.

LVM1 differs mostly from LVM2 by the kernel module, LVM1 uses his own kernel module to manage all storage devices while LVM2 uses the generic device mapper module.

With the new LVM, LVM2, a lot of limits that the original LVM implementation had were either removed or heightened to absurd limits. For example, the original LVM implementation had a maximum Physical device size limit of 2TB. LVM2 heightened this limit to 2EB. While work on the 2.5 linux kernel begun, there were 2 competing implementations, LVM2 and EVMS, made by IBM. With the 2.6 kernel, only LVM2 got in and EVMS remained an out-of-tree project.

This LVM2 implementation based on the Device mapper is the LVM we have now.

Maak uw bijdrage over op banknummer 59.27.19.189 ten name van Stichting Universiteitsfonds Twente.

Op onze website www.utwente.nl/ufonds kunt u makkelijk en veilig via IDEAL een bedrag overmaken.
Daar vindt u ook meer informatie over notariële schenkingen.

Hartelijk dank namens de studenten van de Universiteit Twente.

Met het Universiteitsfonds Twente komen ze verder.

Word nu donateur!

Stichting Universiteitsfonds Twente

De Stichting Universiteitsfonds Twente is een door de Belastingdienst officieel erkend goed doel.
De Stichting heeft de status van Algemeen Nut Beogende Instelling (ANBI).

EMPOWERING the FUTURE

Author: Frits Kuipers

On the15th of May, 2013, Scintilla organizes a symposium themed `Empowering the Future'. The symposium will be held at Spiegel 1 and starts at 8.30 AM. During this day technical presentations will be given, and there will be a demonstration market. If you want to see the latest developments in science and technology, be there!

The current development of civilization shows incredible leaps in science and technology. We try to reach the limit of the human ability to shape and harness the world around us. But this will not happen without our own effort to attain that limit. We are all working towards empowering the human race to realize our maximum potential for mankind's future existence. The possibilities are endless, limited only by the power of the human mind.

> "The possibilities are endless, limited only by the power of the human mind."

The day will start with a nice cup of coffee and an appetizer. After a small introduction from Bram Nauta, chairman of the symposium, the first speaker begin. The symposium consists of seven presentations given by speakers from a wide variety of companies. Amongst others there will be speakers from serveral companies listed on the right.

They will tell us how they think their company and field of expertise will Empower the Future. The symposium committee has attempted to keep the speeches as technical as possible and keep marketing speeches at bay. Everyone who attends the symposium will receive a beautiful document folder containing a credit card sized usb-stick and of course a lunch will be prepared for you.

Furthermore there will be a demonstration market. You will have the possibility to walk by the companies personally and speak with the people working on the technologies of tomorrow. This way you can explore options for your future!

The day will end with drinks at the borrel where you can talk with the speakers of the symposium in a relaxed atmosphere.

Prepare to start Empowering the Future with us on the 15th of May, 2013! Don't forget to check our website and enroll [1].

[1] http://www.scintilla.utwente.nl/symposium

ASML
Lithography systems

CERN
Particle physics research, the Large Hadron Collider

Croon
Electrical infrastructure

Dutch Space
The largest space company in the Netherlands

Imotec
Mechatronics

NXP
Semiconductors

Thales
Aerospace and defense technology

Baking Pies - Part 2

Your own home theatre system

Author: Peter Oostewechel

Welcome to part 2 of this series. Last time a small introduction to the Raspberry Pi was given. Today information is given about the most popular use case for the Raspberry Pi, namely a media center solution. Many people use the Raspberry Pi in conjuncture with a central storage system where they store their series and movies. The Raspberry Pi serves in that case as a frontend for their media collection. There are several distribution that all use XBMC as the media centre frontend. In this article we will discuss one.

Background

The CPU of Raspberry Pi is quite low-end and it can't really decode video at decent speed (even SD MPEG2 is too much). Hardware acceleration has to be used but by default only H264 can be hardware accelerated. You can buy a licence key to enable additional hardware decoders from RPi foundation (currently you can buy MPEG2 and VC-1 license).

The graphics chip on Raspberry Pi is VERY powerful and if video is encoded with a format supported by the hardware decoder, it can easily play HD content in 1080p. But you will have to transcode all your material that is in different a different format.

Hardware decoding only works with the dedicated video player called Omxplayer. This video player is specifically designed for the Raspberry PI's GPU by the XBMC project. It is used in all the different distribution to play HD video. Normal graphical environment (X server) does unfortunately not use accelerated graphics, what causes slow GUI rendering. Luckily XBMC is using Open-

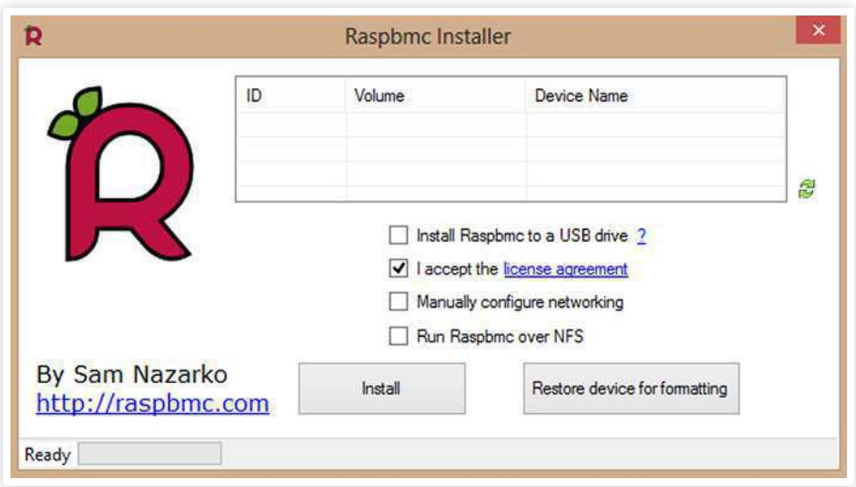

GL ES which does hardware acceleration so it's not going to be a problem.

Installation

For this tutorial we will use the Raspbmc distribution. The Raspbmc distribution is created and maintained by Sam Nazarko and you can download the latest version at www.raspbmc.com. The windows installer will be used in this tutorial but on the site you can also find OSX and Linux installers.

Before launching the application, plug-in the SD card reader and unplug any other USB storage devices to make sure you will not format an incorrect drive by accident.

Now, launch the Raspbmc installer application and select the SD card reader from the devices list. It is recommended to use a USB stick as a storage partition on the Raspberry Pi, so tick the box on "Install Raspbmc to a USB drive". This is because of the speed of most USB sticks. So make sure you buy a fast SD-Card or use a fast USB drive to store the settings and media library on. If you do not use an USB stick, un-tick the box.

After clicking the "Install" button, it will take less than a minute while the installer

downloads and restores the image to the SD card.

After plugging in the memory card, USB stick, USB keyboard, Ethernet cable and power adapter, the computer will boot up and the installer program will run for about 15-20 minutes. The installer even jokes you can get a cup of coffee in the meantime.

Raspbmc will boot up XBMC once the installer script is ready.

In action

You now have your very own ultra-cheap media center. Start by adding media files from your NAS, server or external USB disk. You can also install plugins to automatically download subtitles and cover-art.
One downside to this setup is that the interface sometimes feels sluggish or slow loading fanart. To speed up the interface you can use an USB stick as mentioned earlier or apply the following performance enhancing options.

Overclock it

You can safely speed up the RPi by overclocking it. Go to "Programs → Raspbmc Settings → System Configuration. Next, change the "System Performance Profile" to Fast or Super to overclock the system.

For me, the "Super mode" works reliably when running the Raspbmc off a USB stick. But then I have also installed a small heatsink on the Broadcom chip. There are some know problems with data corruption when using the super overclock mode, running with USB helps to avoid corrupting the SD card that can happen when overclocking.

> "The installer even jokes you can get a cup of coffee in the meantime."

Get a faster memory card

In this case faster is better. Make sure you buy a class 10 memory card to speed things up nicely.

Change the default skin

Confluence skin works fine, but you can improve the responsiveness of the system by using even more lightweight skin.

Quartz skin is very responsive skin, so you could try installing that.

Go to System → Appearance → Skin → Get More... → Quartz → Install to switch your default skin to Quartz.

Conclusion

All in all, the Raspberry Pi starts to be ready to replace your home theatre PC, but it will require some tweaking before you can realize its full potential as a XBMC front-end. The development around Raspbmc looks very promising and after fixing all the bugs and making the operating system more stable, it will be an excellent choice as a very low cost media center.

In the next part of this series the GP-I/O possibilities of the Raspberry Pi will be discussed, so keep in touch to get some idea's to integrate your Pi in your awesome electronic projects.

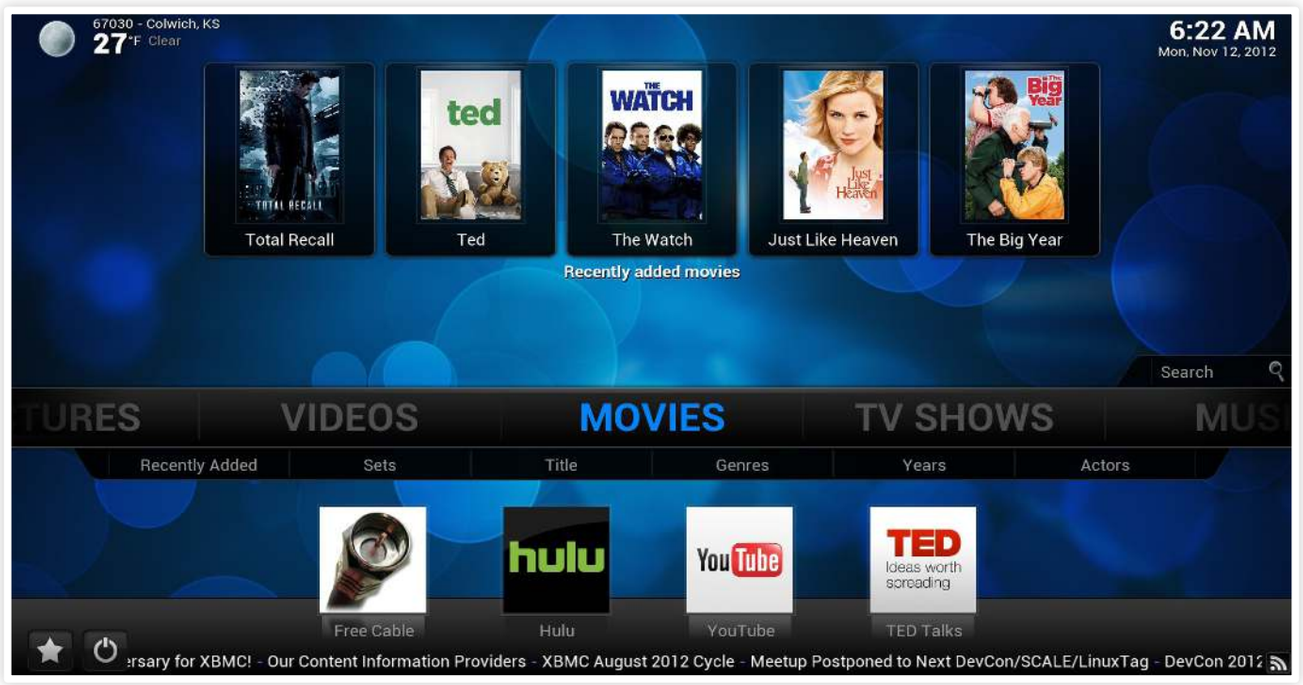

Master thesis

Design and implementation of an analog-to-time-to-digital converter

Author: Dirk-Jan van den Broek

Consider a runner, about to improve his personal record. How do you measure his speed? You could let him run for a fixed amount of time, then shout "stop!" and use a 400 meter long measurement tape to measure how far he got. You could also let him run a 400 meter lap and take his time at the finish line. You could say there is a distance-based measurement method, and a time-based measurement method. Since making an accurate, 400 meter long measurement tape is quite impractical, the choice in this case would be clear: just put an expensive stopwatch at the start/finish line and get it over with. Similarly, in analog-to-digital converters, there are roughly voltage-based measurement methods and time-based measurement methods. However, the choice is not so obvious here. During my master's thesis, I investigated a new, time-based method for analog-to-digital conversion.

The assignment was to design an analog-to-digital converter that takes a two-step approach:
1) convert the input signal to a time difference between two events,
2) convert this time difference to the digital output code.
These 'events' are usually transitions of a digital signal. The idea is shown in figure 1: the input voltage essentially controls a delay that is applied to a reference signal. The time difference td between the reference transition and the delayed transition is a new measure for the input. It is fed to a 'stopwatch' that measures the time difference and outputs the digital data. The controllable delay block will be called the voltage-to-time converter or VTC. The 'stopwatch' will be called the time-to-digital converter or TDC. The reference edge will be called the 'start' signal, and the delayed edge the 'stop' signal.

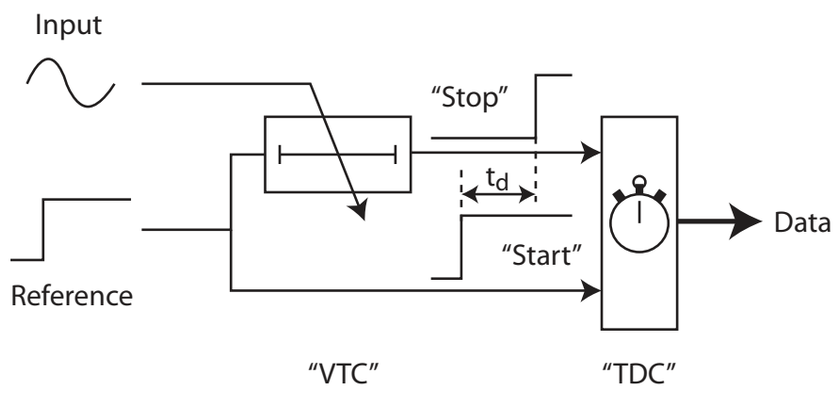

Figure 1: Overview of the analog-to-time-to-digital converter concept.

Motivation

Why would you want to follow this approach for A/D conversion? As most of the readers will know, CMOS chip technology is constantly being improved for digital circuits. Therefore, transistors become ever faster and smaller. Meanwhile the available supply voltages are reduced to around 1 volt in the most modern processes. This makes it more difficult to build analog blocks such as amplifiers, voltage references and comparators, which are key elements in voltage-based ADCs. Meanwhile, digital blocks such as inverters and logic gates become faster and more power efficient. One could say that the 'timing resolution' in modern CMOS has become better than the 'voltage resolution' [1], although this statement is

hard to put into numbers.

Now, the TDC can be a very 'digital' block. In its simplest form it can even be a simple clocked counter, but more sophisticated solutions exist: more on this will follow. So, if the new time-based ADC can already compete with conventional ADCs in current CMOS technology, and the TDC part automatically gets much better with new technology because it is digital, it could be a very attractive solution for the future.

The VTC remains an analog component, so one challenge is to design a VTC that can also withstand technology scaling. Also, the VTC should not dominate the performance of the ADC: if the VTC dominates the area or power consumption of the ADC, then the fact that the TDC improves with technology is meaningless.

Design

TDC

As mentioned before, the TDC can be as simple as a digital counter. However, a counter has a limited clock speed, which limits how many bits can be converted each second. To measure time intervals more precisely, more sophisticated TDCs can be found in literature. In fact, very precise devices were already in use for nuclear research in the 1950's by using vacuum tubes [2]!

The most basic TDC that achieves better time resolution than a counter is shown in figure 2 and is called a flash TDC, because of its analogy to the flash ADC in the voltage domain. It consists of a delay line (for instance inverters or buffers) and a bank of flip-flops. The 'start' signal is fed into the delay line and starts propagating through the delay elements. The 'stop' signal toggles the flip-flops, which take a snapshot of the state of the delay line. The more time between the two transitions, the more flip-flops will capture a '1'.

Many more sophisticated TDCs can be found, such as two-step TDCs, pipelined TDCs, successive approximation TDCs, noise-shaping TDCs: there are almost as many concepts for time-to-digital conversion as there are for analog-to-digital conversion!

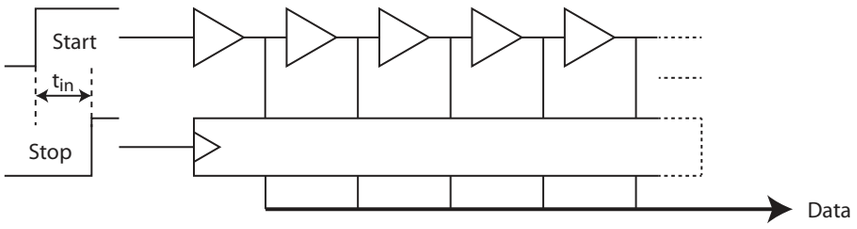

Figure 2: Simplified illustration of a flash TDC.

I decided to keep the design simple, and only address the main drawbacks of the flash TDC. One drawback is that the amount of delay elements and flip-flops quickly becomes impractically large: for example, for 10 bits of resolution, $2^{10} = 1024$ delays and flip-flops are required.

Do we want these high resolutions anyway? Yes, because at lower resolutions this concept has a lot of competition from other fast, low-resolution ADC concepts, for instance the VCO-based ADCs. By targeting 10 bits, we are in a field where the main competition are pipelined ADCs.

> "Do we want these high resolutions anyway?"

To maintain a small area, a ring oscillator was used instead of a delay line: this way, only a few delay stages are needed, plus a counter to keep track of the number of cycles.

Now we have a ring, but where are the edges inserted? The answer is: nowhere. The ring + counter will be continuously running, and therefore the ring becomes a simple ring oscillator. The clue of the design is that the TDC does not 'accept' the start signal, but instead it generates the start signal itself when the VTC has taken an input sample, and is ready to generate the delay. Compare this with the runner standing at the start of a race: The runner determines how fast he will run his lap, but the person holding the stopwatch determines when he will start. The runner is the VTC, the person holding the stopwatch is the TDC.

VTC

In essence, the VTC is a variable delay, which is controlled by the input voltage in a linear way. There are several ways to generate a delay on chip, but as far as I know, they all boil down to one principle: charge a capacitor with a current until the voltage across the capacitor reaches a threshold[1]. The length of such a delay is defined by the magnitude of the current, the size of the capacitor, the initial voltage across the capacitor and the threshold voltage.

Which of these four variables do we control using the input voltage? Controlling the capacitance can be done, but you need voltage-dependent capacitors, which are usually not very linear. Controlling the current requires a voltage-to-current converter, which is also hard to linearize. What remains are the initial voltage and the threshold voltage, which are both very linear ways of controlling the delay. I chose to let the input control the initial voltage on the capacitor, as shown in figure 3. This way, the capacitor has two functions: it is used as a sampling capacitor, and re-used to generate the delay.

Integration of TDC and VTC

An overview of the entire design is shown in figure 4. It is interesting to note one more thing: Now that the 'time base' of the converter is continuously running, we can have multiple VTC's making use of the same TDC at the same time. This divides the

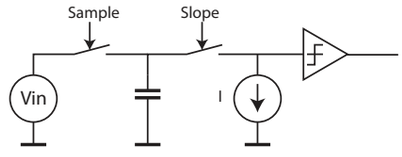

Figure 3: Simplified illustration of a start-voltage-controlled VTC.

[1] OK, you could also charge a coil until the current reaches a certain threshold, but coils are impractical components on chip...

power consumption of the time base over multiple conversions, which helps to keep the power consumption of the ADC low.

Implementation

The majority of the blocks were implemented in 140 nm CMOS. During implementation, some design choices had to be fixed. I chose to implement a 16-stage oscillator, because when the number of stages is a power of 2, the oscillator output can be easily converted to binary code and appended to the counter value. Also, I chose to implement four 9-bit VTC's, but in such a way that these could also be connected together to operate as one 10-bit VTC.

TDC

Ring oscillators with an even number of stages generally do not oscillate. Therefore, differential stages had to be used. A nice paper on a front-end for ultra-wideband radio [3] mentioned a suitable differential oscillator stage, which is shown in figure 5. The flip-flops that record the state of the oscillator are simply existing flip-flops from the digital library.

VTC

To charge the capacitor with a constant current, it would be nice if the current source 'saw' a constant voltage on its output. To accomplish this, the VTC samples the input voltage on the capacitor, and then connects the charged capacitor to a single transistor. This transistor acts as an operational transconductance amplifier or OTA, keeping a constant voltage for the current source. The principle is shown in figure 6.

Reference conversion

In the gain from voltage to output code of the converter, there are several unknowns, such as the magnitude of the current in the VTC, the size of the capacitor in the VTC and the oscillation frequency of the ring. Also, there are low-frequency variations in the gain of the converter, because of temperature differences and supply voltage variations. To compensate for this, the converter will alternatingly convert the input voltage and a full-scale reference voltage. The resulting two values are divided by each other to obtain the actual output value.

> "The converter can indeed consume about 10 times less power"

Results

Figure 7 shows some simulated transients of the system in the 10-bit mode. The system can be seen sampling the full-scale reference voltage, then lifting the capacitor to the correct potential for the VTC and finally converting the input.

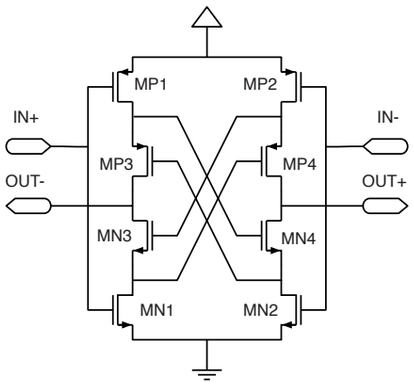

Figure 5: The topology of the differential oscillator stage [3].

An important performance figure for ADCs is the Walden Figure-of-Merit (FoM). It defines how much power the converter uses for each conversion step. If the Walden FoM is calculated for this converter, we obtain 2.32 pJ/conversion step. This is about 10 times too high to compete with state-of-the-art ADCs, however. The power consumption is dominated by digital circuitry. These blocks are expected to improve a lot when the concept is implemented in a newer technology. Initial experiments showed that the converter can indeed consume about 10 times less power when implementing it in 65 nm CMOS instead of 140 nm.

Conclusion

During the thesis I explored a new time-based ADC concept and was able to implement it to a great extent. In CMOS 140 nm it cannot compete with state-of-the-art converters, but initial simulations show that in newer technology it may be an option for the future.

I want to thank NXP for the offer to perform this assignment for them within the ICD group, and for the pleasant cooperation. If you are interested in a bachelor's or master's thesis at ICD or an internship at NXP, I encourage you to contact either the ICD secretary or me directly.

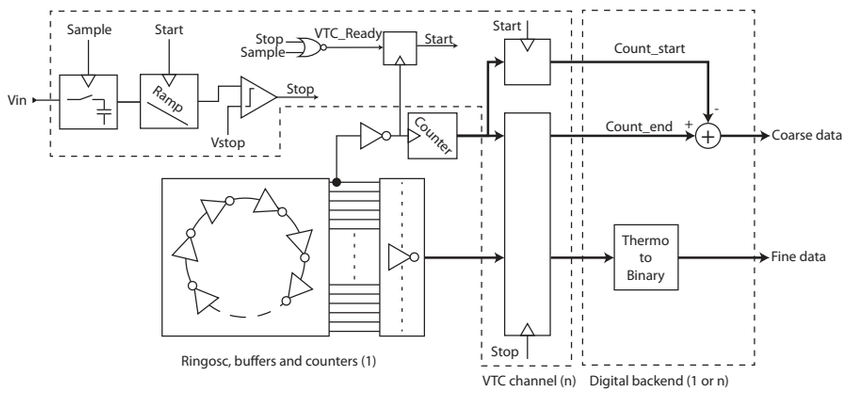

Figure 4: Architectural overview of the proposed system. On the bottom left, the time base of the converter is shown (ring oscillator and output buffers). Within the L-shaped dashed section, one VTC is shown, consisting of a sampling block, a ramp generator, a comparator and flip-flops to record the start and end states of the time base. The other dashed section shows the digital back-end that converts the output values into valid data.

[1] Robert B. Staszewski. Digital deep-submicron CMOS frequency synthesis for RF wireless applications. PhD thesis, University of Dallas, 2002.
[2] Harlan W. Lefevre and James T. Russell. Vernier chronotron. Rev Sci Instrum, 30(3):159–166, 1959.
[3] N. Van Helleputte and G. Gielen. A 70 pj/pulse analog front-end in 130 nm cmos for uwb impulse radio
receivers. IEEE Journal of Solid-State Circuits, 44(7):1862–1871, 2009.

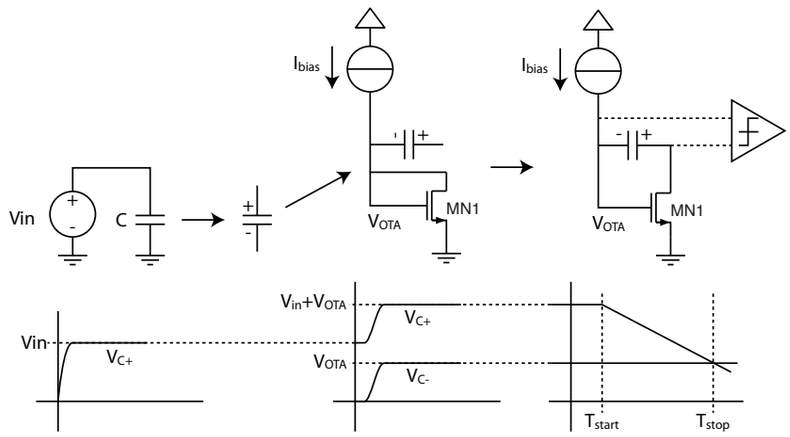

Figure 6: Operation principle of the VTC. The sample capacitor is lifted on top of a diode-connected transistor after the sample is taken. The transistor is then reconnected across the capacitor and starts acting as a virtual ground for the current source, in order to generate a very linear ramp.

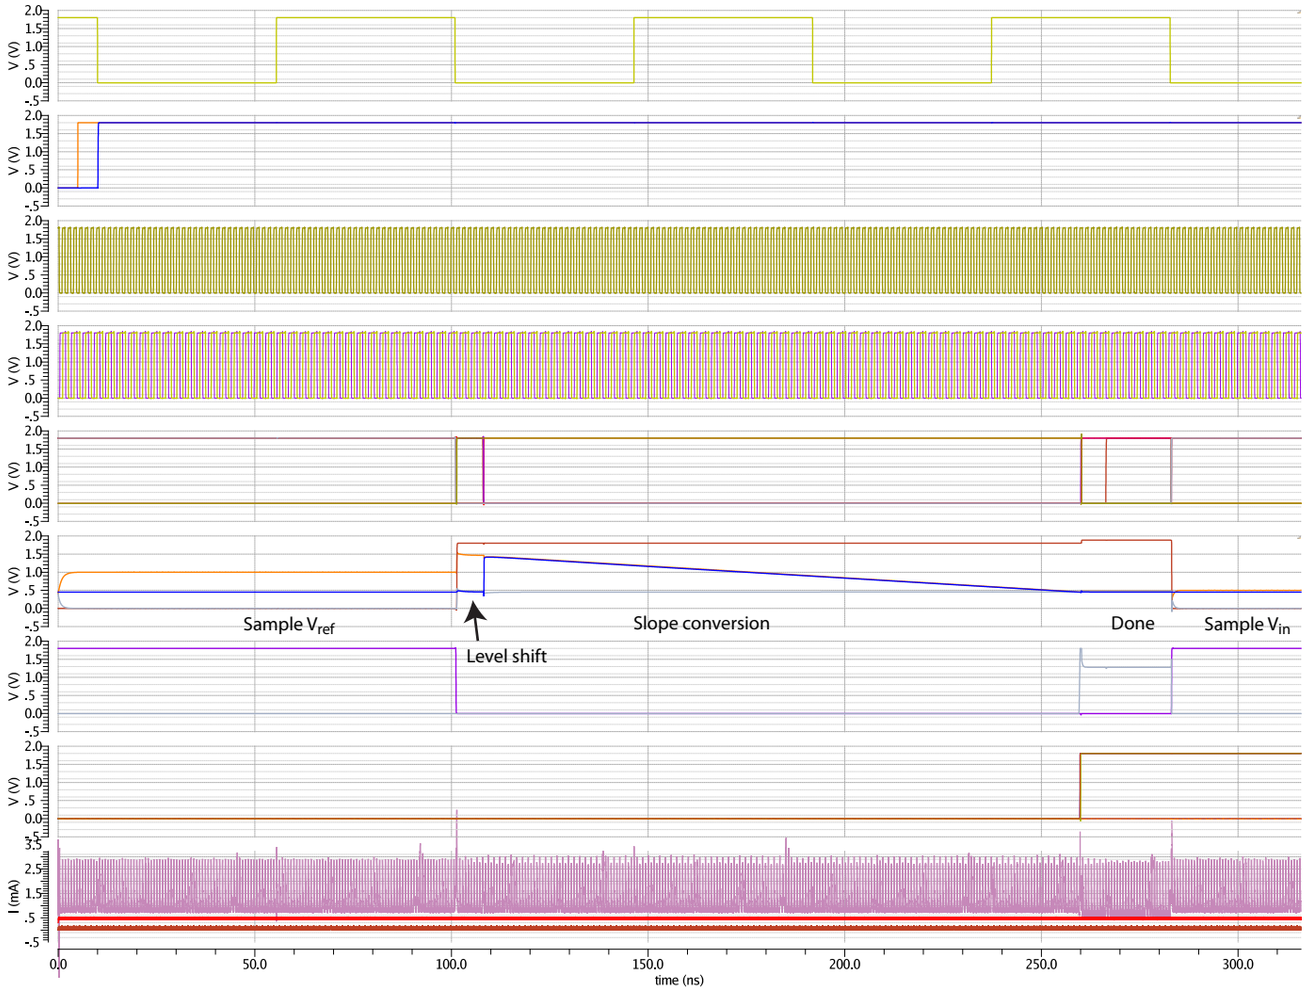

Figure 7: Transient simulation of the complete system in 10-bit mode. From top to bottom: the sample clock signal; a reset signal for the VTC; one phase of oscillator output; the LSBs of a counter; several digital control signals inside the VTC; the analog voltages inside the VTC; the 'start' signal (active low) and the comparator output ('stop' signal); the coarse and fine bit signals settling at digital values at the stop signal; the current consumption of the complete system (purple), oscillator (red) and output buffers (orange).

FRAMES



Leading in Oil & Gas Technology

Frames defines itself by implementing new technologies in a dynamic market with growing internationalization. In addition, we highly appreciate a good working atmosphere in an informal setting.

Well educated ambitious people are working at Frames. Our terms of employment are correspondingly. Frames tries to stimulate its employees to continue learning, both "on the job" as well as through providing training. Our study costs arrangement allows employees to follow an adequate education in an affordable and easy manner.

Specialized in the design, manufacturing, supply, installation and commissioning of complete systems for oil & gas treatment, separation, heat exchanging, flow control and safeguarding, we constantly expand our technological capabilities as well as our customer base to include virtually all the major oil & gas production companies worldwide.

Opportunities to grow in one of our offices are plentiful. Besides developing opportunities in the Netherlands, we are continuously looking for employees who are willing to help with the growth of our international offices.

In total, Frames has five offices in the Netherlands located in Alphen aan den Rijn, Vollenhove, Woerden and Zoeterwoude. Besides expanding in the Netherlands, more and more Frames offices are opened abroad. Supported by offices in Brazil, Germany, India, Malaysia, Russia, Saudi Arabia, UAE and the USA, Frames has managed to make a name for itself in the global oil & gas industry in the past three decades.

The following technology and product groups are part of the Frames portfolio:

- Flow control & safeguarding
- Gas & liquid treatment
- Separation technologies
- Gas processing
- Biogas processing
- Heat transfer
- Automation
- Field services

Interested?
Please visit www.jobsatframes.com or www.frames-group.com for more information.

Frames is a medium sized international company providing turn-key technology for the upstream oil & gas industry. Headquartered in Zoeterwoude the Netherlands with offices in Alphen a/d Rijn, Woerden, Germany, United Arab Emirates, USA, India, Malaysia, Brazil and Saudi Arabia. Our business plan includes a sustainable growth for the coming years. The Frames organization is characterized by an open, international, committed and relatively young staff. Visit www.frames-group.com to learn more about our company.

Vonk Forge

Author: Vulcanus

For this edition the editorial team was looking for something new, something interesting to do. We took a look at our magazine and decided that we wanted to do something with the name "Vonk", which translates as Spark. One can think of a lot of things when thinking about sparks. The most inspiring idea for us was the spark of inspiration. How a simple suggestion or phrase can lead to great work done by ingenious people. This resulted in this new item called Vonk Forge.

In this item we will describe challenges pertaining to electrical engineering, in the hope that this spark will trigger a flash of inspiration in you, the reader. These challenges will be described in detail in every Vonk Forge. One can think of simple electronic projects meant to impress or even be useful. You can then send in your creations to the Vonk. The different projects will then be evaluated by us. And the best projects will be mentioned in the next Vonk Forge. The very best project will even net the creator a special Vonk Forge trophy.

Forge a Vonk

For our first Forging we will start close to home. The challenge will be something meant to impress. We have noticed that during constitution drinks we use a very nice metal staff, it could however be even more awesome. It could shoot sparks when the staff is used to call for attention.

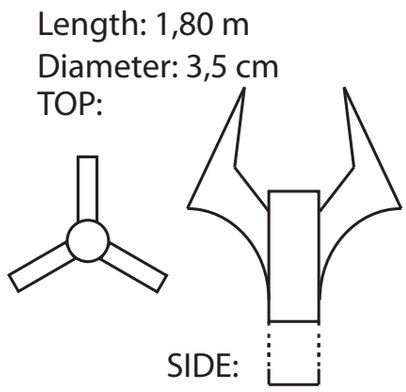

Figure 1: Impression of the staff dimensions

> "The challenge will be something meant to impress."

This brings us to our first challenge. To create a device that will make our staff even better, by shooting sparks from the top, whenever the staff is pounded on the ground. This design should of course satisfy a few requirements. But as this is a challenge of creativity, not all of the requirements have to be met in order to win. However the more of the requirements a project satisfies, the higher the rating for it will be.

The requirements for this projects are simple:

- The device should emit sparks whenever a part of it is pounded onto the ground.
- The device should fit in our staff (as specified in figure 1), barring simple modifications.
- The device should be powered by a battery or another portable way of producing power.

Here are a few tips regarding the device:

- For the sparks: bigger is better.
- As the staff is used during drinks with a lot of background noise, any noise produced is greatly appreciated.
- For the durability we suggest that you keep in mind that the staff will be pounded onto the ground quite forcefully

You can send documentation, including pictures, to my e-mail personal address: vonkforge@scintilla.utwente.nl. You can also use this address to make an appointment to measure, show or test your wonderful contraption.

May the Forge be with you.

Junction

Pepijn Assendorp

Author: Tijmen Hageman

"Nothing is what it seems." Maybe a somewhat cliché sentence, but once again proved correct when observing Pepijn Assendorp. Long hair and a band shirt of "Suicide Silence" indicate a rough personality; something reinforced by the presence of a spiked earring and studded leather gloves. But although this first-year student obviously likes to bang his head to the rhythm of heavy drums, his true nature however is actually rather calm and passive. Find out for yourself on the next few pages!

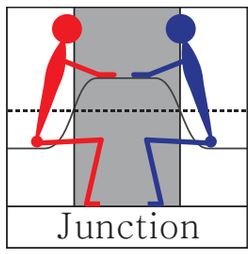

Where are you from?

I'm from Zutphen, where my mother lives. My roots lie in Groningen, but I have also lived in Germany for a while and in Brummen (close to Zutphen). I actually cannot speak German anymore. I could spoke it fluently when I was about four years old, but that has completely gone.

And where do you live at the moment?

Nowadays I live in Enschede, just outside the city center. For the first few weeks I've been commuting, but that wasn't convenient. Now I live in a conventional family home together with three other students. We don't see each other that much, but we always eat together. I enjoy living self-sufficient close to the University.

> "I'm not that agile."

What are your non-study-related activities?

Gaming, and a lot of it. I'm also keen on watching films. I'm not a sports man, as I'm not that agile. Furthermore, I'm active at Scintilla, which I find one of the best things of student life. I'm a member of two committees; StOEL and SOT. I'm a member of StOEL because I liked it and wanted to be able to rely on it myself; I'm a fervent advocate of good education. And I'm a member of SOT because KoenZ assigned me and I did not protest. I haven't got all the experience needed for SOT, but I try to learn and work along. It is great to learn new things you don't learn from the study itself. I've been less active joining Scintilla activities and have only attended the Christmas dinner.

Do you work on EE projects of your own?

No, but If I'd find one I would start on it immediately. It however seems to be compulsory to name your soldering iron, so mine is called Jarvis (from Iron Man).

What is your coolest gadget?

That would be my PC. I've just ordered a new one which will probably arrive next week. It's the first one I've composed and put together myself.

Why have you chosen EE as your study?

It was a given that I would do something technical, that's just the way I am. I started looking at mechanical- and electrical engineering, and I knew I was going to choose either one of them. Eventually, I have chosen EE because I felt like the possibilities are greater when working with electronics.

The best parts of the study are the projects, as you're busy doing practical things. Due to that preference for practical activities, I've doubted whether to attend a higher or academic education. Because of my level I've chosen the latter option. The worst thing about the study so far is programming as that doesn't suit me very well.

> "I'm member of SOT because ... I did not protest"

What are your plans for the future?

If I'm able to keep up, I'll keep on studying EE. It'll be hard, but it's something I want. The future further on is a bit uncertain. On the one hand, I'd like to get to work or reach for a PhD, but on the other hand I'd like to go into teaching. In that case I'll follow a master in didactics in order to acquire a first grade competence in teaching. I've al-

Pepijn Assendorp

Age	Twenty
Study year	First
Birth place	Groningen
Tool	Jarvis

ways been focused on the field of education. Good quality cannot be taken for granted and so I always try to influence the way education is managed.

I do not really have desires to go abroad. I've got all of my family living in the Netherlands and therefore I see no reason to leave. However, if I would get a very good offer, I may reconsider. That is one of the advantages of English education; we're employable across the globe.

Finally, I'm looking forward towards joining Scintilla's board during my third year. I have already considered doing that during my second year, but that would be a bit too soon.

HHDDVVDDBVD

Author: Marcel Wenting

There are many upsides to modern day technology. For the most part it does exactly what machines from the Victorian age also aimed to do, that is make money and make peoples lives easier and more comfortable.

Along with the introduction of such new technologies came trends that only received naming in current times. I'm referring to the terms "innovator", "early adopter", "early majority", "late majority" and eventually "laggards". These terms are coupled to a curve that models the introduction and acceptance process of any new technology.

First are the people with a lot of money and nerds that sell their organs to get it, then come the people that are well informed but didn't have access. Only then comes the first wave of mass adoption. The reason I mention this mechanism is, because it may soon disappear.

You see up to now we engineers have tried our best to make stuff that makes sense to consumers. Take the humble remote control, it comes in many shapes and sizes but because the same icons are used over and over again and the layout is never that different, people use them for a lot of devices.

Sure this did not happen overnight, we need only look at the history of cars to see how difficult the first were to drive. Only after years of iterations the Cadillac type 53 nailed it.

The problem I see now, is that the new technologies are becoming ever more complex and while up to now designers have been able to shield the normal consumer from having to know the inner workings of technology, this may not last.

> "judging by youtube comments, we are headed for the Idiocracy"

Take for example beamers, or any other multimedia device. The function they should perform is relatively simple, but the knowledge needed to actually get them to do what you want is ridiculous.

Couple that to the fact that judging by youtube comments, we are headed for the Idiocracy and you come to the conclusion that soon some people may be excluded from using certain technologies, purely based on the fact that they are unable to do so. This will lead to a "digital segregation" of which the consequences can never be good.

The only hope we have is that with new generations some intuition, some instinct for new technology will develop even stronger than it has done for the past century. For there is no denying that it is remarkable that computers left alone in a village without instruction were being used by curious kids within a week. Kids who had never before heard of or seen such a device, but still they were able to teach themselves how to use them.

Still I feel that be it for economic or moral reasons, engineers will try to keep technology accessible for everyone, but I fear that it is a losing battle.

Puuzle

Author: Truusje

Hello, fellow puzzle maniacs. I hope you have withstood the pressure and cruelness brought on in the first two quartiles. Some of you have fallen bravely, others have lived to fight another day. As usual, I have created a great puzzle to relieve you from the remaining stress.

The slitherlink of the previous edition was successfully solved by several people of whom Diederik van der Valk was the lucky winner. Congratulations to you! My assistants will personally take care of bringing you a delicious cake.

The new puzzle is titled "Light Up" and comes, yet again, from the land of the rising sun. The goal is to light up every square in the grid by drawing light bulbs. The following rules apply:

- Every white square needs to be lit.
- A square is lit if it's in the same row or column as a light bulb, as long as there are no obstructing black blocks.
- No light bulb may light another light bulb.
- A number in a square indicates how many non-diagonally edge squares contain a light bulb.

Geert receives the Vonk-cake for his winning solution of Vonk 30-4

Please send your solutions to truusje@scintilla.utwente.nl or deposit it in the Vonk mailbox in the SK. Good luck and have fun!

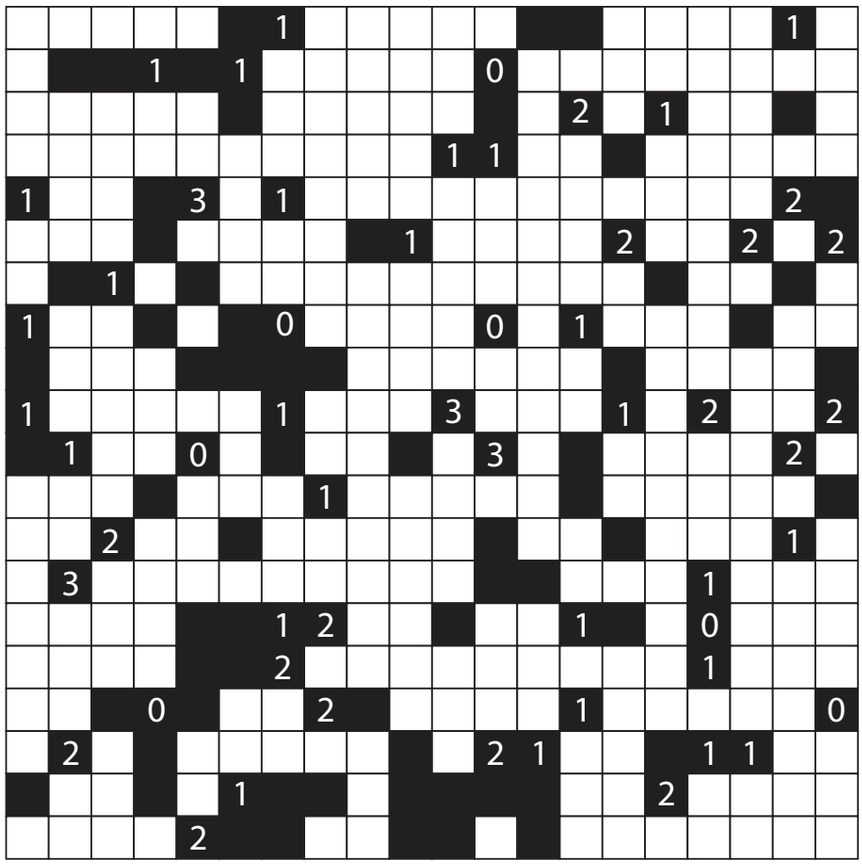

Link yourself to the power of TenneT

Netwerken: daar gaat het om bij TenneT. Letterlijk en figuurlijk. We zijn de eerste grensoverschrijdende elektriciteitstransporteur van Europa met 20.000 kilometer aan hoogspanningsnetwerken in Nederland en Duitsland. Onze focus is gericht op de ontwikkeling van een Noordwest-Europese energiemarkt en de integratie van duurzame energie. Tegelijkertijd staat de continuïteit van de elektriciteitsvoorziening voorop. 24 uur per dag, 7 dagen per week. We zoeken de samenwerking met professionals die interesse hebben in een unieke uitdaging. Wil jij op hoog niveau aan de slag in je vak? Bij een bedrijf dat in meerdere opzichten netwerken verbindt? Link yourself en ga vandaag nog naar

www.werkenbijTenneT.nl

TenneT zoekt:

Ambitieuze technici en andere professionals

TenneT
Taking power further