

Periodiek

Recurring Magazine | Issue 2018-2



Calendar Madness

Gerrit van Tilburg - p. 22



6 - Vector Graphics Primitives

On behalf of the Johan Bernoulli Institute, Pieter Barendrecht and Jiří Kosinka tell us about the mathematics behind vector graphics. How can one create images using Bézier curving instead of pixels? Examples of gradient meshes will be given.




10 - Quantification of Greenhouse Gas Emissions Using UAVs

Huilin Chen and his team used a drone to measure methane emission near a dairy farm. In their drone they used an innovative atmospheric air sampling system called AirCore. Soon this kind of drones will be flying through the skies of Groningen.

26 - Algebraic Chess

How can you make chess even more mathematical? By combining it with linear algebra! Stimulate your brain by figuring out what the solution is to this editions puzzle created by one of the editors, Jonah Stalknecht.



- 4 From the Board
- 6 Vector Graphics Primitives [JBI](#)
- 10 Quantification of Greenhouse Gas Emissions Using UAVs [ESRIG](#)
- 14 Exploding Microdroplets of Tin: The Heartbeat of EUV-light [ZIAM](#)
- 18 Rethinking FMF's Digital Infrastructure
- 22 Calendar Madness
- 26 Brainwork: Algebraic Chess
- 27 Recipe: Greek Rice Salad

From the Editor in Chief

In 1908, the Russian Olympic team arrived at the olympics in London twelve days late because it was still using the Julian Calendar instead of the Gregorian Calendar. In this edition of the Periodiek, you can find out why the calendar was changed and what the possibilities are for alternative calendars.

This is already the third Periodiek of this academic year and therefore it is the third edition that this board of editors has put together. A consequence of this is that the process of creating it goes really smoothly and that we have finished it in record breaking time.

The academic year is coming to an end and a new year will come with new possibilities. One of the possibilities is to help create the Periodiek. Are you interested in joining the board of editors? Don't hesitate to contact us at perio@fmf.nl

-Gerrit van Tilburg

Editors

Jonah Stalknecht,
Gerrit van Tilburg,
Jasper Somsen,
Josselin Kooij,
Rick Vinke

Authors

Pieter Barendrecht,
Jonah Stalknecht,
Thomas van Belle,
Gerrit van Tilburg,
Floris Westerman,
Ronnie Hoekstra,
Josselin Kooij,
Huilin Chen,
Jiří Kosinka

Advertisers

KxA (p.5),
Thales (p.12,13),
ASML (p.20,21),
Belastingdienst (p.25),
Schut (p.28)

Advertise? Contact us at
bestuur@fmf.nl.

Print run 1100 pieces

Press Drukbedrijf.nl

ISSN 1875-4546

The Periodiek

is a magazine from the Fysisch-Mathematische Faculteitsvereniging and appears three times per year. Previous issues can be found at perio.fmf.nl. The board of editors can be reached at perio@fmf.nl.



From the Board

Treasurer

AUTHOR: THOMAS VAN BELLE

Last year I got my Bachelor's degree in Mathematics, and since I did not feel like studying full-time another year, I decided to apply for the 59th board of the FMF. I was hoping to learn some skills that are not thought in either the bachelor or master programme.

Being the treasurer of the association, you have to be somewhat organised. Although my skills seem to be kind of lacking on this point, things are going quite well this year. I think I have learnt a lot since the start of the year. First of all, I found out how the bookkeeping program actually works. This program has a lot of undiscovered functions which mostly are very irrelevant for the FMF. Furthermore, I've learnt to make a list of things I still have to do. I used to just remember everything, but at some point that becomes too much, and forgetting to pay invoices is a bit unprofessional.

Next to learning some useful skills, I found that some big instances are less competent than I expected. Apparently the Belastingtelefoon (a phone service to explain taxes) does not know how taxes work [1], which is kind of the only thing they are supposed to know. Somehow I used to think they were some sort of almighty god of knowledge on taxes. In the end, the FMF's taxes have to be done, so you have to deal with it in some way.



FIGURE 2: Thomas in his FMF suit.

“I have found out how the bookkeeping program actually works.”

Since the work of the treasurer is not visible for the major part of the association, it is a nice change to organise some fun activities. As chairman of the A-team, I organised some awesome activities this year. We had a great hitchhike competition to Bruges, where most teams arrived before dawn, and some nice smaller activities such as swimming and trampoline jumping.

It has been almost a year now since I applied for the board, but it seems like yesterday. As some people say: “Time flies when you are having fun”. So far, it has been a great experience to be in charge of the association with six nice people, and I hope we can continue this for the rest of the year•

[1] <https://nos.nl/artikel/2224499-mysterybellers-teleurgesteld-in-belastingtelefoon-kennis-weer-ondermaats.html>



FIGURE 1: Hitchhiking to Bruges.

Creatieve software engineer?

Hou jij van afwisseling?

Zou je je willen specialiseren in Big data toepassingen?

Geïnteresseerd in de industrie, verkeer of zorg sector?

Ben je goed in programmeren liefst in c++?

Achtergrond in Natuurkunde, Informatica of AI?

Wij hebben vacatures voor zowel ervaren engineers als trainees.

KxA datasolutions is een innovatief bedrijf voortgekomen uit de astronomie, gespecialiseerd in data toepassingen en gevestigd op de grens van Groningen en Friesland. We werken aan veel verschillende dataprojecten. De ene keer houden we ons bezig met slimme datatoepassingen voor het verkeer de andere keer zijn we bezig met de optimalisatie van een fabriek of helpen we mee aan de ontwikkeling van een innovatieve windmolen of analyseren we onverwachte uitval in een fabriek.

Geïnteresseerd? Kom dan eens kennismaken.

Geen auto? Wij kunnen zorgen voor vervoer.



Geïnteresseerd? Bel of mail!

{kxa}
DataSolutions

Wilma Mulder
mulder@kxa.nl
tel. 06 15347819

Vector Graphics Primitives

From Fonts to Gradient Meshes

AUTHORS: PIETER BARENDRECHT AND JIŘÍ KOSINKA

Raster graphics, storing a specific colour for each pixel either directly or through some form of compression, are ubiquitous. The majority of them start their life as a snapshot recorded using a digital camera, nowadays most likely integrated in your smartphone. Though most never leave their medium, some of them find their way to the Internet. It is often there that they meet an entirely different species: the vector graphics.

The alternative representation of graphics

Although designs in vector graphics format are digital entities, they are not defined using pixels. Instead, they are composed using a set of elementary building blocks commonly referred to as *primitives*. Each primitive uses its own mathematical description, which is completely independent of the resolution of your screen. As such, vector graphics are *scalable*: you can zoom in on them without losing quality. An additional advantage is that editing vector graphics is very user-friendly. This is because they are also *stored* as a list of primitives, each with their own geometric or visual properties, and can therefore always be loaded and edited directly.

The number of commonly used primitives is limited. Most of them are curve-based, such as polygons and ellipses. When these are used to define an enclosed

region, the interior can be assigned a colour, or a gradient of various colours.

In this article we take a closer look at the curves used in vector graphics, and then shift our focus to the most expressive primitive available, the *gradient mesh*. We encourage the reader to open a vector graphics editor (such as ADOBE ILLUSTRATOR, INKSCAPE or CORELDRAW) and experiment with the primitives introduced in order to get a better grasp of this way of representing graphics.

Your daily dose of Bézier curves

Many vector graphics primitives are based on *Bézier curves*. In fact, you look at these curves in rendered or printed form on a daily basis! Indeed, the *glyphs* comprising a font are commonly designed using Béziers. Let's take a closer look at the mathematical definition of these curves:

$$C(t) = \sum_{k=0}^d P_k B_k^d(t) \quad (1)$$

with $t \in [0, 1]$, points $P_k \in \mathbb{R}^2$ and $B_k^d(t)$ the Bernstein polynomials of degree d .

Now, this appears rather straightforward, though not immediately intuitive. Why would the sum of the products of two-dimensional *control points* and basis functions produce something sensible? The theoretical




FIGURE 1: The de Casteljau algorithm in action for a quadratic (left) and a cubic curve (right) at $t = 1/2$.




FIGURE 2: Capital G from Linux Libertine and Japanese Hiragana 'ゆ' from Ume P Mincho. The cubic Bézier control polygons are shown in red.




FIGURE 3: Various examples of ligatures using the fonts FreeSans, Linux Libertine and Tengwar Annatar.

That fish **SLUTER!**



explanation is that the Bernstein functions form a partition of unity, so that at every point the curve $C(t)$ is an affine (in this case even a convex) combination of the control points.

Alternatively, the *de Casteljau* algorithm provides us with a geometrical insight. To start, connect the control points P_k with straight edges, together forming the *control polygon*. Now, note that each edge between two subsequent control points can be parameterised as $E_n(t) = (1-t)P_k + tP_{k+1}$. Evaluating all edges at the *same value* for t then results in $d-1$ new points. The procedure can be applied recursively, until we are left with a single point, which in fact lies *on* the curve. More precisely, it is the curve evaluated at the chosen value of t . Figure 1 illustrates the algorithm. Note that the first and last control point are interpolated by the curve, whereas the other points merely *pull* on the curve.

In the quadratic case, we obtain $E_0(t) = (1-t)P_0 + tP_1$ and $E_1(t) = (1-t)P_1 + tP_2$. The next step considers $(1-t)E_0(t) + tE_1(t)$, yielding $(1-t)^2P_0 + 2t(1-t)P_1 + t^2P_2$, which indeed matches Equation (1) for $d=2$.

Although Bézier curves of arbitrary degrees can be defined, cubic curves ($d=3$) are arguably the most intuitive to work with. In this case, the two interior control points that are generally *not* interpolated can be assigned the role of *tangent handles* (they determine the tangents of a Bézier curve at the first and last point). This provides the designer with sufficient freedom to control the shape of the curve. However, there is only so much that can be modelled with a single curve. As such, the common workflow is to connect multiple curves to design more intricate shapes. Positioning the tangent handles on either side of a shared control

point such that the three are co-linear then yields a tangent-continuous transition.

Both quadratic and cubic Bézier curves are used in font design (for TrueType and OpenType fonts, respectively). Naturally, quadratic curves can be expressed as cubic curves (which is, after all, the desired format in many vector graphics editors). In the power basis $\{1, x, x^2, \dots\}$ this is trivial, but when using the Bernstein basis, it is slightly less so. Still, the control points to represent a quadratic curve as a cubic one can be readily obtained through a procedure known as *degree-elevation*. Figure 2 shows the control polygons of two glyphs, where the tangent handles have been omitted.

There is a lot more that can be said about fonts. Directly related to designing the outline of individual glyphs is the readability of a font. Note that this is a rather relative notion. For example, most people probably prefer a *monospaced* font for writing code (`LATEX`, C++, ...). On the other hand, on your personal blog you might want to use something more aesthetic. Concepts such as *Kerning* (the spacing between consecutive characters) and *ligatures* (combining multiple glyphs into a single glyph) become important and have to be considered in the design process. Figure 3 shows examples of ligatures.

To conclude this section, we mention that there is a plethora of other spline types. However, only a few are actively used in vector graphics. Among those are *B-spline* curves, which can be interpreted as a more efficient way of representing multiple Bézier curves which connect with automatic continuity. In addition, there are *spiro splines*, which lend themselves well for more organic art, though they are also used for the design of fonts.

Going bivariate in full colour

Now that we are more familiar with curves, we can also construct *curve networks*. For the moment, we will constrain our expressiveness and only consider regular quadrilateral meshes (not unlike the the wall of tiles in your kitchen). Assuming Bézier representation, we assign a colour to every control point (not to the tangent handles). Every edge can now interpolate between the colours assigned to its endpoints. Immediately, this raises the questions of which colour space to use (e.g. sRGB, HSV or CIELUV) and which blending functions to apply (e.g. linear or cubic) to interpolate. Assuming an RGB space for now, we can focus on the blending functions. A natural approach is to extend our existing Bézier curve used for describing the geometry (the curve network) to also cover our colour transitions. As the control points already have colours assigned to them, this amounts to assigning colours (or colour differences) to the tangent handles.

Next, we would like to fill the enclosed regions (*patches*) defined by the curve network with a visually pleasing colour transition. Linear and radial gradients might be available [BB13], but are typically not versatile enough. We thus decide to use some form of *transfinite interpolation* to define a colour gradient in the interior of our patches. To do so, we can again use a Bézier representation, now in tensor-product form. It might however be more practical to adopt a tensor-product Hermite description, which is a change of basis compared to Bézier. We obtain a 5D control structure (2D geometry and 3D colour), which now also has interior data. In Hermite form, we can set these interior components to zero by default (after all, we are not provided with information about the interior). Technically, we now have Ferguson patches,

a subclass of Hermite patches. Both can also be seen as specific types of Coons patches, a term that is used in some vector graphics editors.

This concise description covers the basic approach taken by existing implementations of gradient meshes (although some of them might use slightly different blending functions).

Upon using this primitive for a while, it becomes clear that it comes with certain limitations. Most noticeably, the topology of the network is rather restrictive. In order to design organic objects with a regular quad mesh, we might have to stretch and compress patches considerably. However, a more severe shortcoming of the fixed topology shows when we want to refine the mesh (e.g. in order to express local detail). Refinement in this context means the insertion of additional control points, which in this case implies the insertion of points along an entire row or column of the mesh. It goes without saying that introducing control points where they are not needed can only be counter-productive, either by slowing down the software, or cluttering the workspace and complicating the designer's job. As such, some form of *local refinement* would be an attractive addition. Although the proposed feature is conceptually straightforward, the implementation of a proper user interface and data structure are more involved [BLHK18]. Figure 4 compares the existing approach of *global refinement* to our local refinement.

Another improvement regarding the topological structure can be made by allowing what we refer to as *branching*. Essentially, this means that the curve network can be composed of branches (each a regular

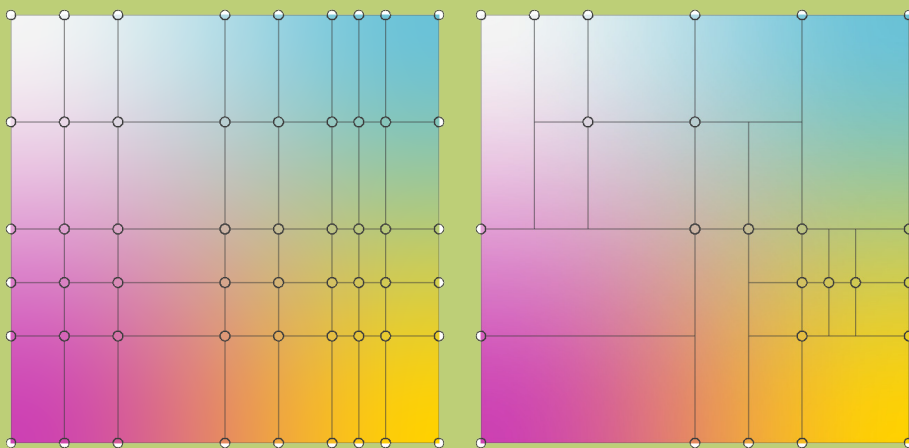


FIGURE 4:
Global refinement
(left) versus our
local refinement
(right) in gradient
meshes.

quad mesh) without introducing extraordinary points in the interior of the network (that is, control points with a valence not equal to four). Figure 5 shows an example taking advantage of both topological relaxations.

Real-life objects often have sharp colour transitions. Unfortunately this is not quite supported by the traditional gradient mesh primitive, unless degenerate patches are used or multiple gradient meshes are placed directly next to each other. Taking inspiration from *diffusion curves* [OBW+08], an alternative approach to work with flexible colour gradients in a vector graphics context, we allow the assignment of colour to individual *sectors* of control points. Figure 6 shows two examples using this feature.

Do these extensions completely resolve the restrictions of traditional gradient meshes? Not quite, but at the very least they are a good start in making the primitive more versatile. Looking at ongoing future work, we mention the extension to curve networks of arbitrary topology. The work presented in [LKSD17] proposes a solution, but lacks the possibility for local refinement and has some other shortcomings. In addition, our current local refinement procedure is merely *dyadic* refinement (bisecting the parameter domain). Although this suffices for some designs, refinement of patches at arbitrary parameter values would be a welcoming addition.

The examples of gradient meshes in this article were created manually. A separate line of research considers the automatic creation or conversion from raster graphics, which is referred to as *vectorisation*•

References

[BB13] Pascal Barla and Adrien Bousseau. Gradient art: Creation and vectorization. In *Image and Video-Based Artistic Stylisation*, pages 149-166. Springer, 2013.

[BLHK18] Pieter J. Barendrecht, Martijn Luinstra, Jonathan Hogervorst, and Jiří Kosinka. Locally refinable gradient meshes supporting branching and sharp colour transitions. *The Visual Computer*, 2018. To appear.

[LKSD17] Henrik Lieng, Jiří Kosinka, Jingjing Shen, and Neil A Dodgson. A colour interpolation scheme for topologically unrestricted gradient meshes. In *Computer Graphics Forum*, volume 36, pages 112-121. Wiley Online Library, 2017.

FIGURE 5: A plum modelled using our extended gradient mesh primitive, supporting local refinement and branching. The curve network is depicted on the left.





FIGURE 6: Two more extended gradient meshes using our local refinement and sharp colour transition features.



[OBW+08] Alexandrina Orzan, Adrien Bousseau, Holger Winnemöller, Pascal Barla, Joëlle Thollot, and David Salesin. Diffusion curves: A vector representation for smooth-shaded images. *ACM Trans. Graph.*, 27(3):92:1-92:8, August 2008.

Quantification of Greenhouse Gas Emissions Using UAVs

AUTHOR: HUILIN CHEN

The atmospheric mole fractions of major long-lived greenhouse gases (GHGs), i.e. carbon dioxide (CO_2), methane (CH_4), are increasing due to human activities. The rise of the greenhouse gases in the atmosphere is the main cause of global warming. The 2015 United Nations Climate Change Conference in Paris (COP21) agreed to limit global warming “well below” 2°C and, if possible, below 1.5°C . This requires not only massive reductions or even negative emissions in greenhouse gases, but also advanced understanding of their sources and objective monitoring of the emissions reduction efforts.

Available atmospheric measurements of greenhouse gases have been employed to determine their sources and sinks through so-called inverse modeling; however, these estimates are associated with large uncertainties, mainly because of sparse observations and transport errors. Apparently, more long-term observations would help reduce the uncertainties in the estimates although they are expensive. Large international efforts are being made to monitor the atmosphere, e.g. surface and tower measurements within the European research infrastructure Integrated Carbon Observing System (ICOS) and the NOAA ESRL Global Greenhouse Gas Reference Network, and remote sensing measurements from TCCON, and from satellites such as GOSAT, OCO2, and Sentinel-5 Precursor. Here I present my work on quantifying greenhouse gas emissions using unmanned aerial vehicles (UAVs), which is an emerging research line that is complementary to ongoing research activities.


“Large international efforts are being made to monitor atmospheric greenhouse gases.”

Why a UAV?

For atmospheric research, aircraft is one of the most important platforms to perform large-scale vertical and horizontal measurements. However, aircraft is not suitable for measurements below ~ 150 m due to strict aviation regulations and for the sake of safety. To this end, an unmanned aerial vehicle, commonly known as a drone, has its great advantage to fill the vertical gap of measurements between the surface and the lowest altitude usually reachable by aircraft. Besides this, a UAV is also cost effective.



FIGURE 1: Sample analysis near a dairy farm.



FIGURES 2 and 3: The measurement device and the team with the UAV flying nearby.

UAV AirCore for greenhouse gas measurements

AirCore is an innovative atmospheric air sampling system invented by Dr. Pieter Tans at NOAA. It consists of long, thin-wall stainless-steel tubing capable of sampling and preserving atmospheric profile information. AirCore is evacuated as it is lifted up to a high altitude (~30 km) by a balloon, and then during descent after the balloon bursts it is passively filled with atmospheric air samples due to the increasing ambient pressure. The samples are analysed on the ground to retrieve the GHG vertical profiles.

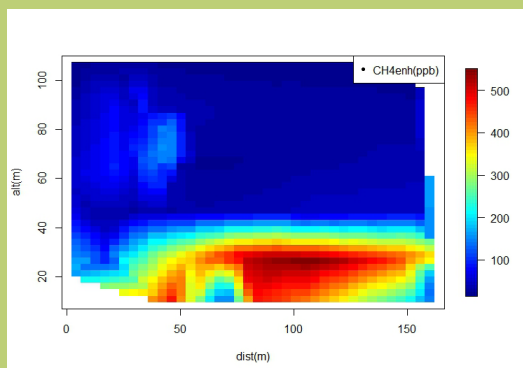
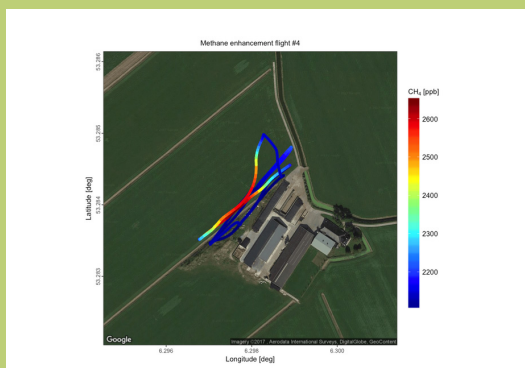
Based on the concept of AirCore, we developed a UAV-based active AirCore for atmospheric mole fraction measurements of CO₂, CH₄, and CO. As opposed to the conventional concept of passively sampling air using the atmospheric pressure gradient during descent, the active AirCore collects atmospheric air samples using a pump to pull the air through the tube during flight, which opens up the possibility to sample atmospheric air in both vertical and horizontal planes for a wide variety of opportunities, see one example below.

Quantifying CH₄ emissions from dairy cows on a farm in the Netherlands

According to the emissions inventory, dairy cows are the largest emitter of methane in the Netherlands. However, the uncertainty in the estimate may be large, as no large-scale direct measurements have been made. Here we aim to use the UAV AirCore to accurately quantify anthropogenic emissions of CH₄ from dairy cows on a farm.

We flew the UAV AirCore downwind a farm at various heights, and then use so-called mass balance approach to quantify the total amount of CH₄ flowing out of the farm. By using the exact number of cows on the farm, we are able to calculate an averaged emission rate of cows from the farm. This work in progress is part of an EU-funded Marie Curie training network.

More recently, I was funded by the United Nations Environmental Program to quantify emissions from the midstream and downstream oil and gas sector in the Groningen province of the Netherlands. We will be soon flying drones often over the skies of Groningen•



FIGURES 4 and 5: The flight path and measurement of methane by the UAV near the dairy farm.



Judie Ibrahim

Student Aviation Studies,
Hogeschool van Amsterdam, 23 jaar

“Toen ik op zoek was naar een stage heb ik veel bedrijven aangeschreven, van drie bedrijven kreeg ik een aanbieding. Ik koos voor Thales omdat het een internationaal georiënteerd bedrijf is met meer dan 85% export. Tijdens mijn opdracht bij Thales hield ik mij bezig met het analyseren van het kostprijscalculatie proces en onderzocht ik waar dat mogelijk te verbeteren. Aangezien ik bij Aviation Studies mij vooral bezig houd met techniek zocht ik juist een stage waar ik mij bezig kon houden met bedrijfskundige processen. Het onderwerp van de opdracht bij Thales was mij onbekend, maar juist daarom wilde ik ermee aan de slag en zag ik het als een uitdaging.

Ik ben van mening dat je met een oprechte interesse theorie wel eigen kunt maken. Toen ik bij Thales aan de slag ging was ik verrast over de open cultuur. Je kunt iedereen aanspreken en er wordt tijd voor je vrijgemaakt om je verder te helpen. Als je zelf initiatief neemt dan liggen er veel kansen. Na mijn advies te hebben geïmplementeerd ben ik van plan om in september te beginnen met een de Master Business Administration aan de UvA in Amsterdam.

Ik denk dat dat de combinatie van een technische Bachelor met een Bedrijfskundige master je helpt om met twee verschillende brillen naar oplossingen te zoeken. Ik onderhoud contact met mijn stagebegeleider en hoop dat ik na mijn master bij Thales aan de slag kan.”

“Als student is je grootste angst dat je afstudeeronderzoek in een lade verdwijnt. Bij Thales word je echt serieus genomen en is er vertrouwen in je. Zo wordt mijn advies omtrent een nieuwe methode van kostencalculatie zelfs in 2016 geïmplementeerd.”

Op zoek naar een stage, afstudeeronderzoek? Start jouw carrière bij Thales. K

“Tijdens mijn studie was ik mij al aan het oriënteren op de arbeidsmarkt. Ik was op zoek naar een technisch bedrijf met een interessant product waar ik mijn natuurkunde achtergrond zou kunnen inzetten. Op de BètaBedrijvenBeurs in Nijmegen kwam ik in gesprek met een recruiter van Thales. Zij nodigde me uit op het hoofdkantoor om samen de mogelijkheden bij Thales te bespreken. Eén van de vacatures sprong er voor mij meteen uit, namelijk die voor de functie van Trial Conductor. Hiernaar heb ik gesolliciteerd en na een paar gesprekken werd ik aangenomen. Inmiddels werk ik al een paar maanden met veel plezier bij Thales.

Als Trial Conductor bouw ik voort op de tijdens mijn studie opgedane kennis. Regelmatig zal ik met een team naar het buitenland gaan om onze radarsystemen op marineschepen te testen. Daar laten we op zee aan de klant zien dat het systeem inderdaad zo goed is als beloofd in het contract. Dit doen we door middel van allerlei scenario's. Zo laten we bijvoorbeeld een F-16 invliegen om te zien wanneer de radar deze voor het eerst detecteert.

Tijdens mijn sollicitatie kreeg ik ook een aanbod van een ander technisch bedrijf, maar de goede sfeer bij Thales was voor mij doorslaggevend. De mensen hier zijn erg behulpzaam en nemen de tijd om je iets uit te leggen. Thales voldoet zeker aan het beeld dat ik ervan had: een high tech bedrijf met een fijne werksfeer.”

“Als Trial Conductor bouw ik voort op de tijdens mijn studie opgedane kennis. Regelmatig zal ik met een team naar het buitenland gaan om onze radarsystemen op marineschepen te testen.”



Annelot Schuring

Afgestudeerd in Natuurkunde,

Radboud Universiteit Nijmegen, 24 jaar

studeerplek of eerste baan?
klik op www.thalesgroup.com/nl

Exploding Microdroplets of Tin

The heartbeat of EUV-light sources for nanolithography

AUTHOR: PROF. DR. IR. RONNIE HOEKSTRA

The research efforts of the Quantum Interactions and Structural Dynamics group (Thomas Schlathöler and Ronnie Hoekstra, who is also group leader at the Advanced Research Center for Nanolithography, Amsterdam) focus on the ultrafast processes induced in interactions of energetic particles or photons with matter ranging from single atoms and molecules via clusters to surfaces. We aim at unveiling and linking 1) the primary electron dynamics occurring on femtosecond timescales and 2) the subsequent structural and functional changes. The knowledge obtained is applied in fields such as nanolithography, space research, mass spectrometry and, last but not least, radiotherapy.

In Groningen, for our experiments with ions we operate the ZERNIKE-LEIF facility. This unique facility with its highly-charged ion source, beam lines and auxiliary equipment is uniquely suited to investigate interactions on the atomic scale which are key to novel extreme ultra-violet (EUV) light sources for nanolithography as developed and manufactured by ASML based in Veldhoven, the Netherlands. ASML is the world market leader of lithography tools for application in the semiconductor industry.

Keeping Moore's law alive: EUV sources for nanolithography

Moore's law states that the number of transistors on an affordable CPU doubles every two years. Ever since the seventies, Moore's law is very much alive. This is good news for the Dutch economy, which thrives on its outsized, one fifth share of the global market for semiconductor equipment valued at 40B€ annually by 2020. Keeping up with Moore's law drives the semiconductor industry and enabled innovative and disruptive technologies such as flash memory, which is currently displacing hard drives. The world is very much in need of such novel storage concepts to cope with the tens of zettabytes (10^{21} or 2^{70} bytes) of data that are expected to be generated by 2020.




FIGURE 1: Typical series of sideview shadowgraph images of the expansion of tin droplets exposed to a weak laser pulses of 0.5, 1.5, and 5 mJ. Note how the laser energy influences the details of the expansion and break up.

Keeping Moore's law alive provides an enormous challenge to science and industry. To meet this challenge, sources of EUV light at 13.5 nm wavelength will be used to drive the next generation of nanolithography machines. Its introduction is happening right now, with early adoption expected in coming years by leading-edge chip makers. The step towards EUV is crucial to the continuation of the miniaturisation of the features on chips. Present-day, state-of-the-art "193i" immersion lithography uses 193 nm wavelength light. The shorter the wavelength, the better the resolution, therefore features with much higher imprinting resolution can be printed which enables the continuation of Moore's Law well into the future, shrinking computer CPUs and flash memory to ever smaller sizes. However, the generation of the required EUV light of sufficient power, reliability, and stability provides a continuing challenge.

Highly charged tin ions in plasma of exploding droplets of tin are the atomic sources of EUV light at 13.5 nm. Given the substantial industrial and societal interest in plasma EUV sources, it is remarkable how little is actually known about the physical processes that generate EUV light. In part this is due to the special characteristics of the EUV generating tin plasma which lie in terra incognita between profoundly studied solar and lighting plasma

environments.

In Groningen we do basic research on the interaction of tin ions with matter on the atomic scale. The research is performed in direct collaboration with the Advanced Research Center for Nanolithography (Science Park Amsterdam) which is founded (2014) in response to the need of ASML for basic research with a long-term application perspective and a potential for breakthrough solutions in the generation and use of extreme ultraviolet light for lithography. The Advanced Research Center for Nanolithography (ARCNL) is a public-private partnership between the Netherlands Organisation for Scientific Research (NWO), the University of Amsterdam, the Vrije Universiteit Amsterdam, and the semiconductor equipment manufacturer ASML.

Laser-produced plasma sources of EUV light.

In a laser produced plasma (LPP) source of EUV light a several 10 kHz stream of tin droplets is hit by high-power infra-red laser pulses. The produced plasma emits the wished for EUV light which is captured and refocused using a hemispherical, multi-layer "collector" mirror.

In an industrial EUV source, droplets of molten tin several 10 µm in diameter serve as mass-limited targets. They are first irradiated by a relatively low-




FIGURE 1: Overview of ASML's extreme ultra-violet (EUV) lithography tool for the semiconductor industry.




FIGURE 2: Schematic overview of different types of plasma, based on temperature and density.

energy laser pulse, varying from 10 ps to several 10 ns in pulse length, to mould these droplets into optimal target shapes. This “prepulse” is still so strong (10^9 - 10^{11} W/cm²), to generate a dense ablation plasma which is responsible for propulsion of the target droplet and its subsequent deformation into a “pancake” shape. After a certain time delay the deformed droplet is irradiated with the higher intensity (10^{10} - 10^{11} W/cm²) high-energy main pulse. This pulse ablates the droplet and creates a high-density plasma at relatively high temperatures of roughly 100 eV (~1 MK).

The average charge state Z of the tin ions lies near $Z=10$. Line emission from these electron-impact excited highly charged tin ions Sn^{8+} - Sn^{14+} , exhibiting a serendipitous alignment of atomic transitions from $4p^6 4d^{m-14f}$, $4p^6 4d^{m-15p}$, and $4p^5 4d^{m+1}$ states towards the $4p^6 4d^m$ ground states ($m=0$ -9) largely independent of the charge state. No other element in the periodic table of elements exhibits such a strong overlap of wavelength of the strongest transitions when changing its charge state. Since within a plasma species occur in a wide range of charge state the fact that wavelengths in tin ions coincide near 13.5 nm is an enormous bonus of using tin plasma to generate

EUV light and is the origin of the high conversion efficiency of drive laser power into 13.5 nm EUV light.

Why does one use 13.5 nm and not for example 10 or 15 nm? For EUV light which needs to be transported in vacuum not to be absorbed, hardly any optics is available. The optics’ elements are basically restricted to multilayer mirrors. Within the EUV range Mo-Si mirrors are the most efficient ones with their reflectivity peaking in a bandwidth of 2% around 13.5 nm. It is the mirrors that have define the optimal EUV wavelength for the sources and it is the EUV generating plasma that needs to be tuned to the mirror’s avelength.

Besides the useful EUV light, the plasma also produces unwanted out-of-band radiation (i.e. emission outside the 2% bandwidth around 13.5 nm), as well as fast neutral and ionic “debris”. Besides limiting efficiency, out-of-band radiation is detrimental to source performance. To make the source as efficient as possible, and to extend its lifetime, it is crucial to understand the atomic processes that contribute to the emission processes at the fundamental level. The fundamental process here is the line emission by highly charged tin ions in the plasma. The extreme density

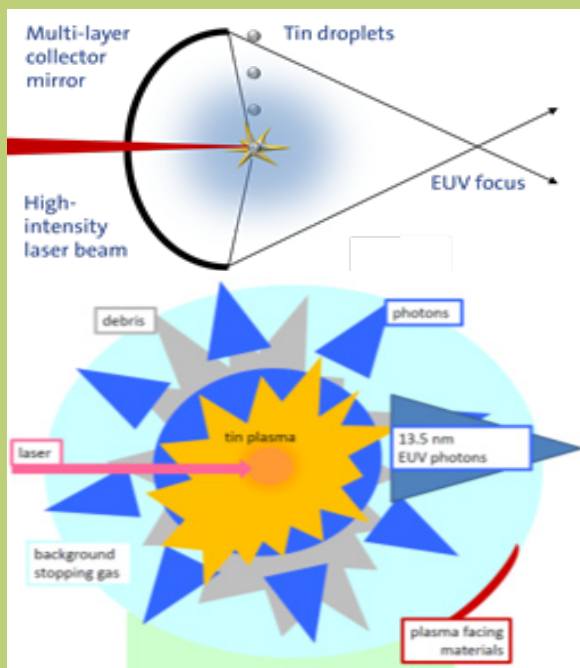


FIGURE 3: Top: Schematics of a droplet-based LPP EUV source providing 13.5nm light for nanolithography. A stream of tin droplets is hit by high-power laser pulses. The produced plasma emits EUV radiation which is captured and refocused using a multi-layer “collector” mirror. Bottom: Artistic view of the ingredients of the exploding tin plasma and its surroundings.

understand the creation of the massive particles. The H_2 buffer gas is also helpful in reducing the speed of the ions ejected from the plasma. On the other hand, EUV light might be absorbed by the gas, an unwanted effect. To find out optimal settings of the gas pressure requires high-quality input data on in the interactions between tin ions, neutrals and clusters with H_2 , data that is by and large fully missing.

Source physics aspects thus range from atomic, molecular, and optical physics, through plasma physics all the way to the physics of fluids. The physics questions posed above relate to three prime missions towards the best EUV sources: maximise light conversion efficiency, minimise the impact of fast ions, and limit neutral debris. Answering the questions will help erect the scientific scaffolding that will underpin achieving the industrial goal of building the most efficient, high-intensity EUV plasma sources for nanolithography. Recently taken into routine operation, the ZERNIKE-LEIF ion facility is ideally suited to unravel the interactions of multiply charged tin ions with buffer gas and mirror surfaces at the most fundamental level. Most relevant ion energies and charge states to be investigated will be determined on the LPP source installed at ARCNL•

of the plasma makes the full understanding of the atomic plasma processes a very rich challenge by itself and moreover most of the fundamental interaction data is non-existent. The complex, many-electron structure of the relevant tin ions hampers accurate theoretical calculations of the relevant cross sections. Therefore it is experimental data that is called for.

Besides photons, a LPP EUV source also produces fast ions, which are ejected from the plasma of the exploding tin droplets. These ions can travel with enormous speeds and strike nearby expensive EUV optics like bullets. This bullet barrage is extremely detrimental to the lifetime of the source in general and the ~€1 million multi-layer mirror in particular. Understanding how these ions are created allows one to figure out how to tune the experimental parameters in such a way to minimise their impact.

A third product is neutral debris, a term encompassing atoms, clusters, nanoparticles, and microdroplets. Atoms are easily removed from the source by means of a buffer gas, but the more massive particles do reach the plasma-facing surfaces and coat them with an ever increasing film of tin. Therefore, it is crucial to

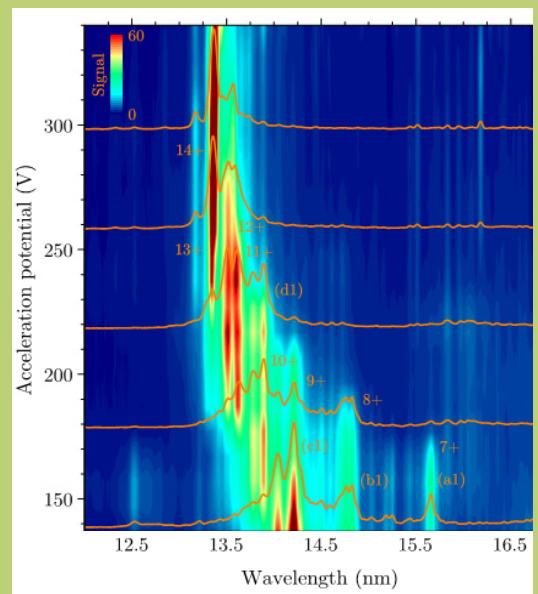


FIGURE 3: Alignment around 13.5 of EUV light emission by Sn ions. Experiments were performed at the EBIT facility at the Max Planck Institute in Heidelberg. By increasing the voltage on the EBIT (vertical axis) the dominant charge state of the trapped Sn ions increases from 9+ to 14+.

Rethinking FMF's Digital Infrastructure

An explanation by the Comcie

AUTHOR: FLORIS WESTERMAN

At first sight it might not be so clear, but the FMF has quite a lot of digital infrastructure in place to support the association. The most obvious ones are the website, the mail and the computers in the NSFW and the Vector Space. Less obvious, but nevertheless important, are things like the member registration system and the streepsysteem. All these things are set up and maintained by us, the Comcie.

Single point of failure

You might know us mostly from our infamous Facebook posts that there'll be some maintenance to the server(s), which then takes more time than expected. Luckily, most of the time things run just fine and there's not much to do other than answering the incidental question or resetting a password. However, when shit hits the fan it used to hit quite hard. This became particularly clear a year ago, when we suffered from weeks of downtime following an incident.

This particular incident made us completely rethink the setup of our digital infrastructure, in order to prevent such a massive outage in the future. The old setup had one major flaw: (almost) everything ran on a single machine, our trusty server Babylon. It ran the mail server, the web server, the LDAP server (for managing user accounts), the file server and (obviously) a Minecraft server all at once, in the same operating system.

“It ran the mail server, the web server, the LDAP server, the file server and a Minecraft server all at once.”

This formed a single point of failure: one problem and all services would be offline. The setup was also very brittle: everything had a dependence on everything and even the smallest change would cause everything to collapse. In short: the situation was far from ideal and after the incident it was clear we had to get our shit together and do it properly.

Virtualisation

The first step we took was to virtualise all our systems. For those not familiar with virtualisation: you basically run multiple (virtual) computers at the same time on the same (physical) machine. This allowed us to separate all the services; the mail would run on a different virtual machine (VM) than the website for example. So when the mail would malfunction, a reboot of that VM would have no effect on the website. This makes the entire setup a lot more versatile, flexible and reliable.

To allow for virtualisation, you first need a so-called hypervisor: a piece of software that coordinates distributing the computer resources (RAM, CPU time, etc.) to the VMs. For this, we chose to use Proxmox which is actually an entire operating system based on Debian. Since it is mainly geared towards enterprise use, it is very stable and brings many enterprise features as well, such as the ability to combine multiple physical machines into a cluster that can seamlessly distribute running the VMs to properly balance the load across all machines.

At the moment we have 5 VMs, all running Debian. They are responsible for the LDAP server, the file server, the mail server and the website(s). For the last category we have two different VMs, we will discuss the details later on. For storage, we use two 8TB drives in a ZFS mirror configuration, with a 120GB SSD for caching and journaling. This makes sure that the loss of one hard drive does not mean that all our data is lost, and the SSD adds some speed to the entire system; it is basically not noticeable that we're 'only' running on hard drives in 2018.

Websites

Besides all the services listed before, the FMF has multiple websites. A lot, to be exact. Before we started setting up the new system, only the main website (fmf.nl) ran on Babylon, most others still ran on our old server, Atlas (RIP 15-01-2018). The main website now runs in its own dedicated VM, with no problems ever since.

The other websites did pose a problem to us, because the code was mostly very old (databases that were created in 2001, anyone?) and required old versions of PHP or even Perl to run, introducing many security risks and limiting the use of new techniques. We would need to have multiple VMs running, each with their own base of Debian taking up approximately 1GB of memory. That is of course not really desirable or flexible.

To solve this problem we make use of Docker to run every single website in its own container. A container is basically a very lightweight VM: the kernel (the "base" of the operating system) is shared between the host and the container, which saves a lot of loading time and other resources. Therefore, starting one is a matter of seconds and a simple website only takes around 10-20MB of RAM to run. Docker allows your application to be split up in many different parts and for each individual part the amount of system resources available can be adjusted; this makes it even easier to set up a flexible and scalable system.

At the moment, we run every single website (committee-pages like toetscie and fotocie, event pages and former-board sites, to name a few) in a separate container with its own configuration, adapted for the website's needs. By default, Docker is managed in the command-line, which is not the most user-friendly. To solve this, we use Rancher, a web interface that



FIGURE 1: One of the high-tech computers in the members room.

allows us to manage Docker on a higher abstraction level. In the web interface we define the websites that should be run and Rancher will make sure that the service is always available and it will reboot the containers if there are problems automatically.

We also make use of Traefik as a reverse proxy, a system that accepts all incoming requests and forwards them to the right container. Traefik integrates natively into Docker, saving a lot of trouble with IPs and other network-related things. It also automatically requests Let'sEncrypt certificates and handles the HTTPS-traffic, making it easier to start using that without changes to the website code.

All of this together makes the current digital infrastructure of the FMF a lot more flexible, reliable and above all, a lot more secure. This flexibility should make it easier to keep up-to-date in the future and should prevent another massive outage.

The Comcie

Are you terribly confused? Are you interested in learning the details? Come to our borrellecture on June 7th, 17:00, where we'll show you a live demo of how this all works, and answer all your questions! If you're as enthusiastic about this as we are, join us! We're looking for new members who want to keep up the Comcie's pride and like to fiddle around with the latest technology. We do not require you to be completely up-to-date on all these techniques and we're more than willing to teach you!•

ASML: Where technological progress has far-reaching benefits. Be part of it!

These days, a small USB stick costing only €10 can hold up to 16 GB of data. In hospitals, a camera the size of a pill can be swallowed to survey a patient's intestines. Modern pacemakers, critical devices that control abnormal heart rhythms, are now less than a tenth the size of earlier ones. And in the oceans, tiny GPS transmitters track endangered turtles to help protect them.

While these devices are incredibly small, they represent a big milestone in technological progress. At the heart of each of these life-enhancing innovations is a microchip – a tiny package of integrated circuitry that powers the performance of the device.

In a world in which major breakthroughs measure only a few nanometres in size, the constant quest is to produce chips that are smaller, faster, more effective and less expensive. One of the major high-tech players leading the quest is ASML, a manufacturer of lithography systems for producing computer chips.

Crucial step

ASML, located in Veldhoven in The Netherlands, supplies equipment to all the world's major chip manufacturers. These include Samsung, Intel and TSMC.

There are dozens of steps along the path to producing a chip. ASML helps manufacturers take just one of these steps, but it's a very crucial step: lithography. Lithography involves exposing and chemically etching the wafers used to 'print' a chip's components. The more accurate the lithography process is, the smaller the resultant microchip can be.





Using ASML's latest generation of machines, it's possible to print lines on chips that are only about 20 nm thick. To put this into perspective... that's like printing the contents of a 500-page novel onto a centimetre-long strand of human hair!

Opportunities to be part of progress

To continue leading the race to produce smaller, faster, cheaper chips, ASML is looking for people who always want to do better and never give up. People who love what they do – not because it's easy, but because it's hard. People who will help find the best idea, the best solution, the best way forward.

People pushing technology further

The driving force behind ASML's technological breakthroughs is its forward-thinking engineers. ASML's more than 16,000 employees are some of

the most creative thinkers in the world of physics, mathematics, chemistry, mechatronics, optics and informatics. And because ASML invests over €1.0 billion annually into Research & Development, these experts have all the resources at their disposal to push progress to the extreme. It's the only way ASML can maintain its edge – worldwide.

A place of learning

ASML is an ideal environment for professional growth and development. The company offers a fulfilling career, not just a job. ASML rewards employees competitively and provides coaching, training and personal career development. Flexibility, enthusiasm, ambition and customer focus are the foundation for a world of opportunity. To find your opportunity, visit www.workingatasm.com•

ASML

Be part of progress

Calendar Madness

AUTHOR: GERRIT VAN TILBURG

Proposal for a new calendar: 13 equal sized months, 28 days in each month, the first day of each month is always a sunday, the last day is always a saturday. That is far more logical than what we're currently doing!

This is called the international fixed calendar which was proposed in 1902. Of course, everyone believing in the bad luck of friday the 13th may be against this idea, but there must be other reasons why this calendar is not implemented yet. Why is our calendar the way it is, what are the alternatives and what do the alternatives draw down?

13 months like this:						
SUN	MON	TUE	WED	THU	FRI	SAT
1	2	3	4	5	6	7
8	9	10	11	12	13	14
15	16	17	18	19	20	21
22	23	24	25	26	27	28

FIGURE 1: The international fixed calendar.

1	Martius
2	Aprilis
3	Maius
4	Iunius
5	Quintilis
6	Sextilis
7	September
8	October
9	November
10	December

FIGURE 2: The original month names.

History of the calendar

The Egyptians were the ones who determined that one solar year is approximately 365,25 days. The time between two full moons is roughly 29 days and this is the reason why a lot of calendars have 12 months. The calendar which we are using nowadays is obtained after numerous modifications of the Roman calendar, originating from 753 BCE. It is said to be invented by Romulus, the first king of Rome. It consisted of a total of 10 months, with 6 months of 30 days and 4 months of 31 days. The calendar year only lasted 304 days with 61 days unaccounted for in the winter. This calendar did not work as it did not align with

the seasons. This was fixed in 700 BCE by shortening the lengths of the months and adding Ianuarius and Februarius and getting rid of the “winter days”. However, the total number of days in a year was 354 or 355 days, so some extra days were added in the end of the year, called the intercalary month.

When Julius Caesar was in charge, he performed a Caesarean section on the Roman calendar and gave birth to the Julian Calendar. He got rid of this intercalary period and changed the length of the months such that a year consisted of 365 days. He also added the concept of leap years, an extra day once every 4 years. Since January and February were now

the first and second month, the names of the months did not make sense at all. For example: octo represents eight, but October was the tenth month. Julius Caesar and Augustus Caesar saw this as an opportunity to claim the names of the 7th and 8th month, by naming them July and Augustus, respectively. Apparently, after these two guys no one successfully claimed the name of a month. That is why we still use illogical names for the months September till December.

The Julian calendar is the basis of our current calendar, but it had its flaws. Nowadays we know that the length of a year is approximately 365.2524 days, so a leap day once every 4 years is too many. Pope Gregory XIII introduced the Gregorian calendar in 1582, but its implementation took a long time. Turkey was the last country to switch to the Gregorian calendar in 1927. The Gregorian Calendar uses a leap day rule which works better, but is more complicated. Currently (1901–2099), the Julian calendar is only 13 days behind the Gregorian calendar.

Alternative Calendars

The International fixed calendar (IFC)

The international fixed calendar (Figure 1) has 13 months of 28 days and one extra day called new-years day, and in leap years there will be an extra new-year's day. The biggest advantage of this calendar is its consistency. For example, saturday will always be on the 7th, 14th, 21th or 28th of a months. So friday the 27th can be pay-day forever and your salary will come in regular intervals.

Unfortunately, 13 is not divisible by 4 and therefore one can not split up the year in 4 quarters of 3 months. However, a week is still 7 days long and therefore there will be 52 weeks in a year, which is divisible by 4. So the people who are really committed to quarters can just divide the year in 13 weeks.

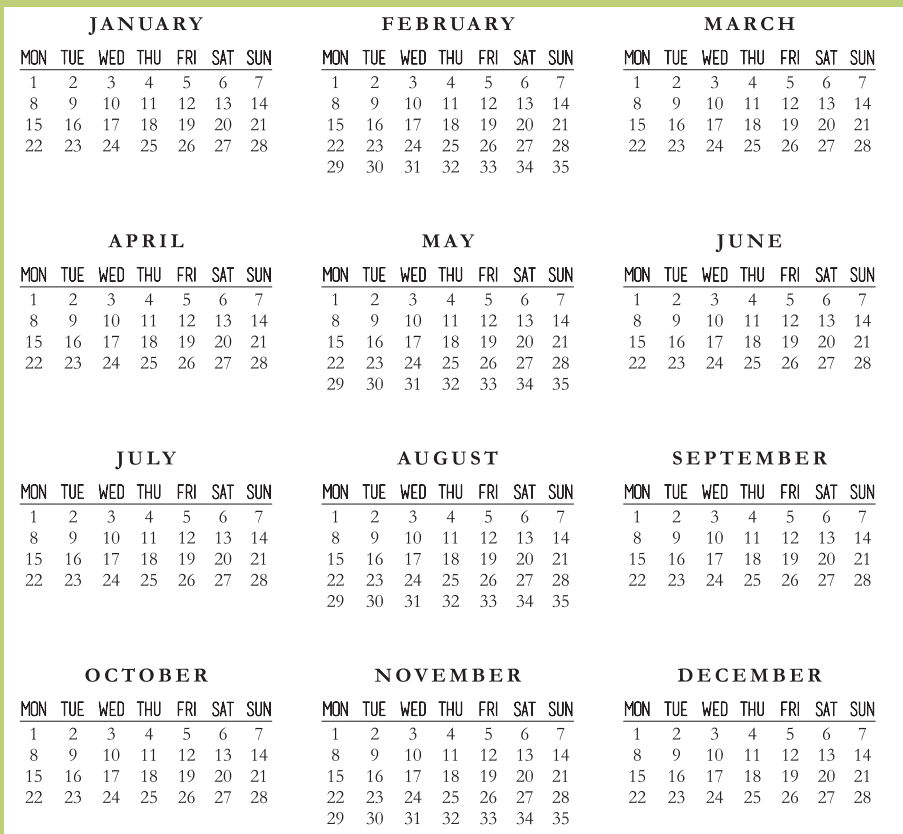


FIGURE 3: The symmetry 454 calendar, with a 4-5-4 week structure.

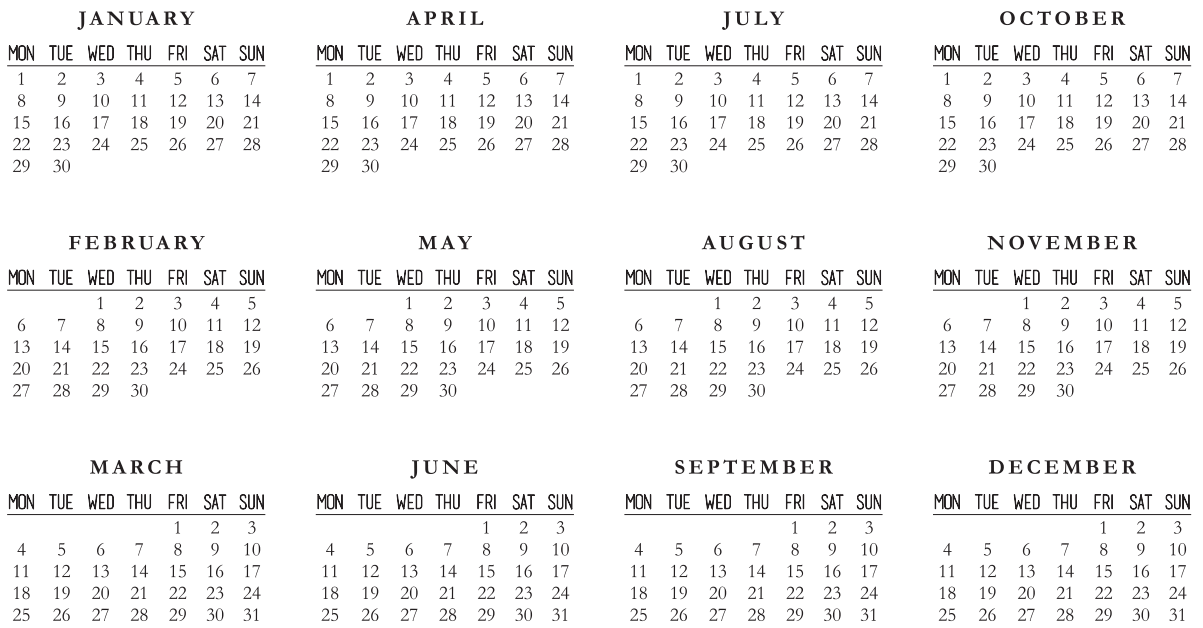


FIGURE 4: The Hanke-Henry Permanent calendar, in which which each quarter has the same structure.

Symmetry454 (SYM454)

The Symmetry454 (Figure 3) Calendar is also a calendar that conserves the traditional 7-day week. It has symmetrical equal quarters and starts every month on Monday. A quarter consists of 3 months which have a length of 28, 35, 28 days. Unlike the International Fixed Calendar, there are no intercalary new years days, but once every 5 or 6 years, a full week is added to december.

Hanke-Henry Permanent Calendar (HHPC)

The Hanke-Henry Permanent Calendar (Figure 4) is quite similar to the Gregorian calendar. It consists of eight 30 day months and every third month is 31 days. Instead of leap days, it has an extra week every 5 or 6 years. Years would proceed regularly as each day would occupy the same position as it had the previous year.

What about birthdays and holidays?

It is preferable that people who are older than others have their birthdays before others in the new calendar as well. This is why most people, except some born in January, get a new birthday. To calculate what your new birthday will be, you need to look at the

number of days into the year up to your birthday. For example, my birthday is on the 14th of October, which is 287 days into the year. This is saturday the 7th of November on the IFC and sunday the 8th of November in SYM454.

Possibility of a calendar revolution

We've come far already, with a calendar that is accepted almost everywhere. Our calendar is broken in its fundamentals: it is not able to address all of the issues associated with measuring and organising the passing of time. It would be more logical to adapt a different calendar, but it seems quite hard to do so.

There are numerous reasons why this would never work, but by looking at the past it actually might be possible to globally adapt a new calendar. Once, it seemed unthinkable that indoor smoking would be forbidden. On the 3rd of september in 1967 Sweden abruptly switched driving lanes from left to right. Ideas that once seemed impossible have been adapted in the past, so why should we not implement an alternative calendar?•



Belastingdienst



*Dit is Maarten. Data Scientist
bij de Belastingdienst.*

*Hij werkt aan slimme data innovaties die
de maatschappij extra miljoenen oplevert.*

Samen met Maarten op basis van data de Belastingdienst optimaliseren?
Ga naar werken.belastingdienst.nl/data-en-analyse.

Brainwork: Algebraic Chess

AUTHOR: JONAH STALKNECHT

Each board represents a legitimate chess position with white to play. On each board white has one unique way to checkmate in one move. White moves up the board on every board.

Rules for multiplying pieces:

- Multiplication is commutative
- White piece \times White piece = White piece
- Black piece \times Black piece = Black piece
- White piece \times Black piece = Empty square
- Anything \times Empty square = Empty square
- The multiplication rules is NOT necessarily the same for Black and White pieces
- For each piece of a colour, there exists a piece of the same colour that together multiply to make any other piece of the same colour.

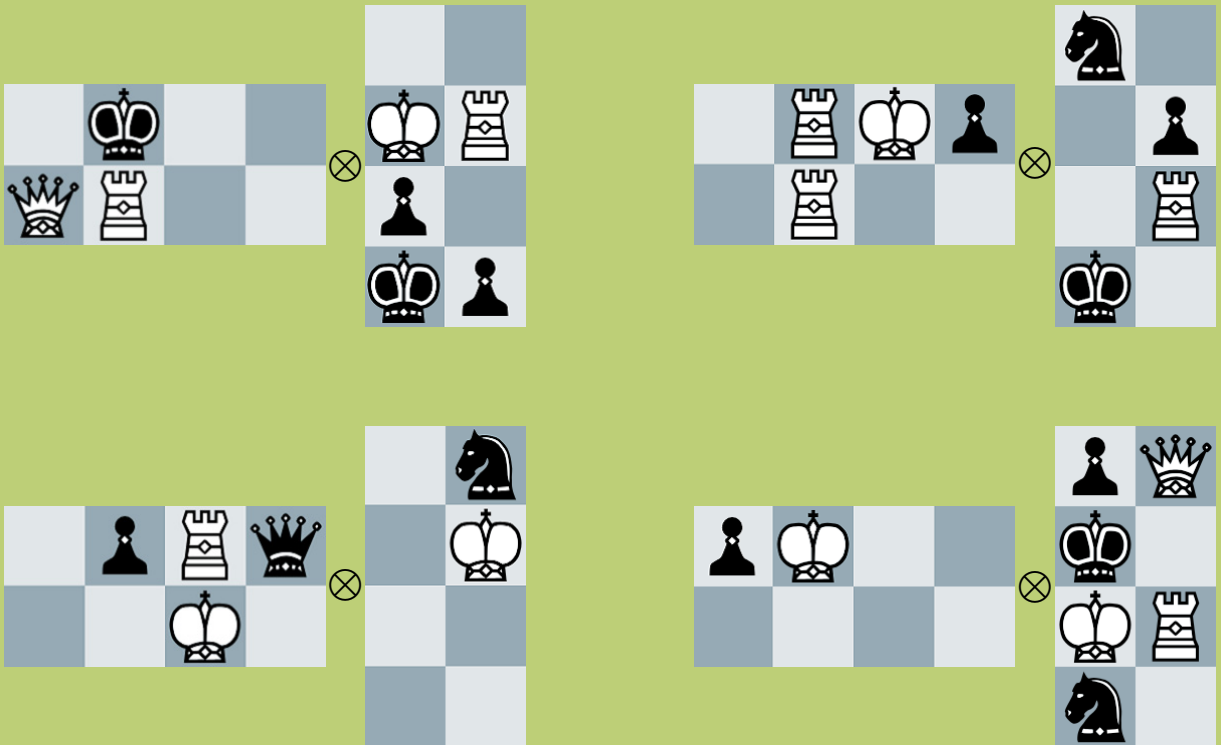
HINT:

$$A \otimes B = \begin{pmatrix} a_{11}B & \cdots & a_{1n}B \\ \vdots & \ddots & \vdots \\ a_{m1}B & \cdots & a_{mn}B \end{pmatrix}$$

Goal:

Determine every board state and find checkmate for white in $(1 \ 1) \otimes \begin{pmatrix} 1 \\ 1 \end{pmatrix}$

You can send your solution to perio@fmf.nl before October 15th•



Previous Brainwork

We've had many enthusiastic (and frustrated) responses to the previous brainwork. If you still didn't finish working on it, don't read any further because the solution will be given at the end of this text. There were different ways to solve it, some figured out a way to factorise out the values for the letters in modulo 10000000, some bruteforced letters, but the online anagram solvers were popular as well.

The people that sent in the solution were (in order of receiving the correct solution): Klaas Hakvoort, Harmjan de Vries, Thomas ten Cate, Bart Dopheide, Thijs Havinga, Eric Jager, Jeroen Lammers, and Johan Jager. They can pick up their (small) prize at the NSFV. Ask a board member or one of the editors.

The solution was the motto of our very own University of Groningen: *Verbum Domini lucerna pedibus nostris.*

Greek Rice Salad

Summer Recipe without Lettuce

This recipe is perfect for a sunny afternoon in the Noorderplantsoen. It takes about 15 minutes to make a big batch and can be eaten cold.

AUTHOR: JOSSELIN KOOLJ

Ingredients (per person)

- 75 grams of rice
- 1/4 cucumber
- 1/4 red onion
- cherry tomatoes
- 1 tsp. basil
- salt and pepper
- 50 grams of feta
- 1/2 paprika
- 2 tbsp. Greek yoghurt

Recipe

Boil the rice as is instructed on the package. Cut the cucumber, onion, tomatoes, feta and paprika in cubes of approx. 1 cm³. Put these in a bowl and add a pinch of salt and pepper the basil. Mix this before adding the rice. Add the Greek yoghurt and mix again. Eat it directly or let it chill in the fridge and eat it cold•





*“Let's go to the
Noorderplantsoen.”*

Tip: add some olives.

Tip: if you like meat, add some smoked chicken.



FIGURE 1: Greek Rice Salad.



Schut Geometrische Meettechniek is een internationale organisatie met vijf vestigingen in Europa en de hoofdvestiging in Groningen. Het bedrijf is ISO 9001 gecertificeerd en gespecialiseerd in de ontwikkeling, productie, verkoop en service van precisie meetinstrumenten en -systemen.

Aangezien we onze activiteiten uitbreiden, zijn we continu op zoek naar enthousiaste medewerkers om ons team te versterken. Als jij wilt werken in een bedrijf dat mensen met ideeën en initiatief waardeert, dan is Schut Geometrische Meettechniek de plaats. De bedrijfsstructuur is overzichtelijk en de sfeer is informeel met een "no nonsense" karakter.

Op onze afdelingen voor de technische verkoop, software support en ontwikkeling van onze 3D meetmachines werken mensen met een academische achtergrond. Hierbij gaat het om functies zoals *Sales Engineer*, *Software Support Engineer*, *Software Developer (C++)*, *Electronics Developer* en *Mechanical Engineer*.

Je bent bij ons van harte welkom voor een oriënterend gesprek of een open sollicitatiegesprek of overleg over de mogelijkheden van een **stage-** of **afstudeerproject**. Wij raken graag in contact met gemotiveerde en talentvolle studenten.

Voor meer informatie kijk op www.Schut.com en Vacatures.Schut.com, of stuur een e-mail naar Sollicitatie@Schut.com.



APPROVE
for DeMeet

