[OBJECTIVE

on innovation and technology

Technolution

NO.12 NOVEMBER 2009

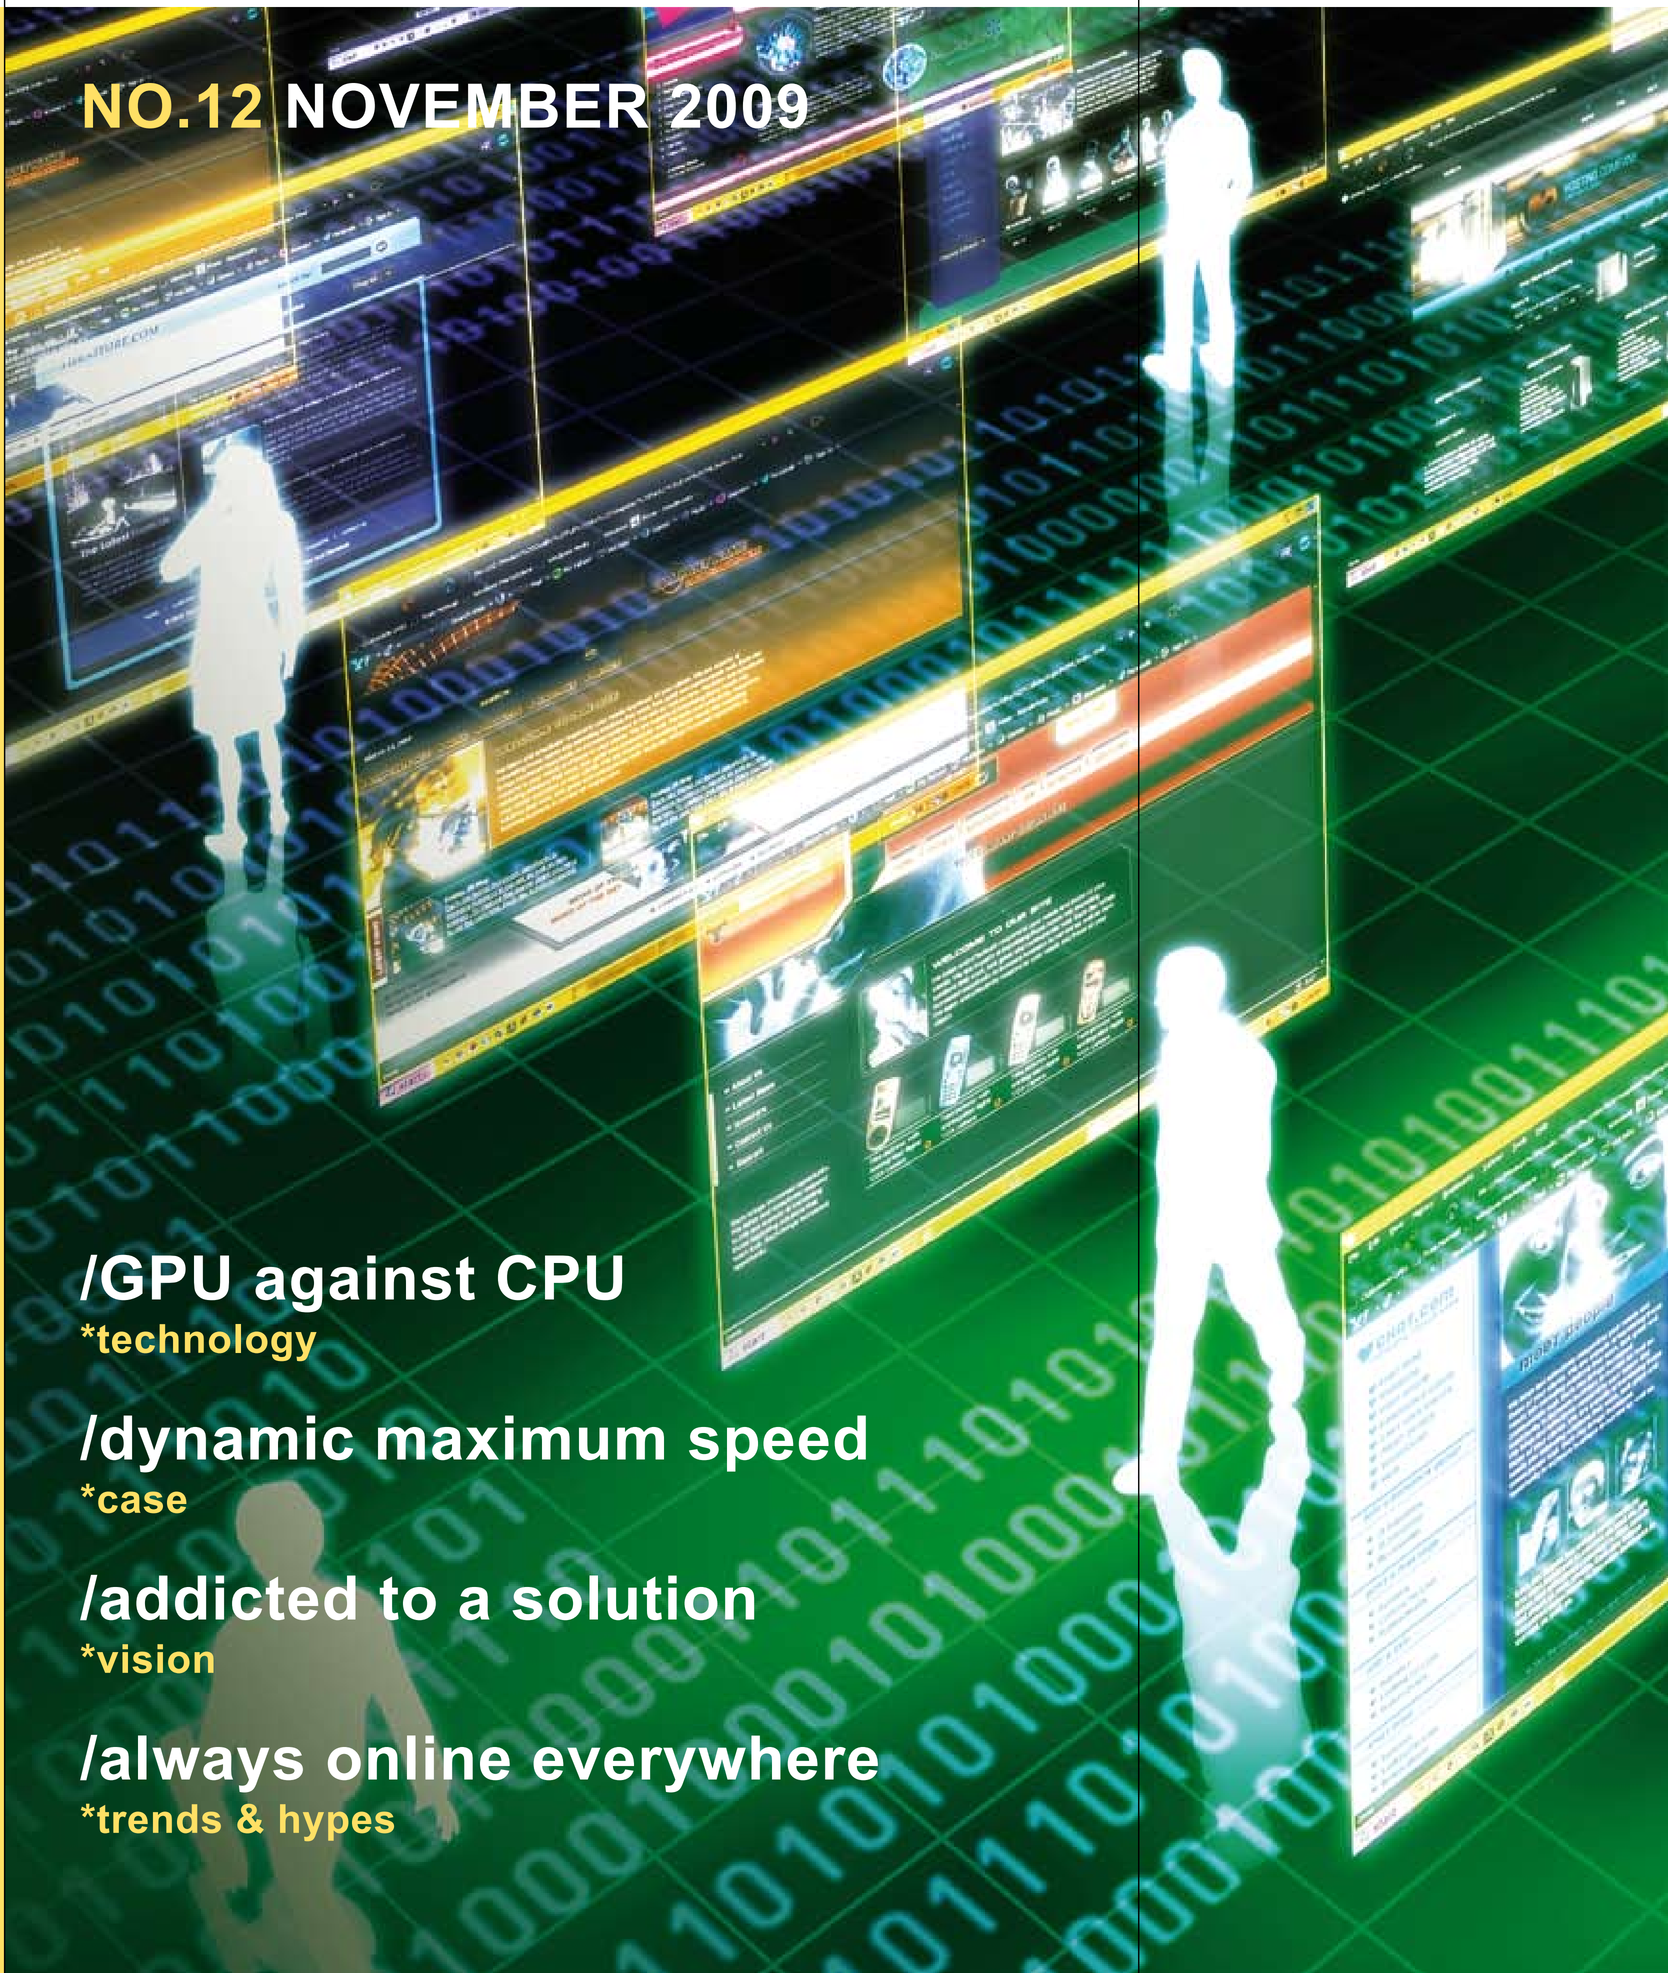

/GPU against CPU
*technology

/dynamic maximum speed
*case

/addicted to a solution
*vision

/always online everywhere
*trends & hypes

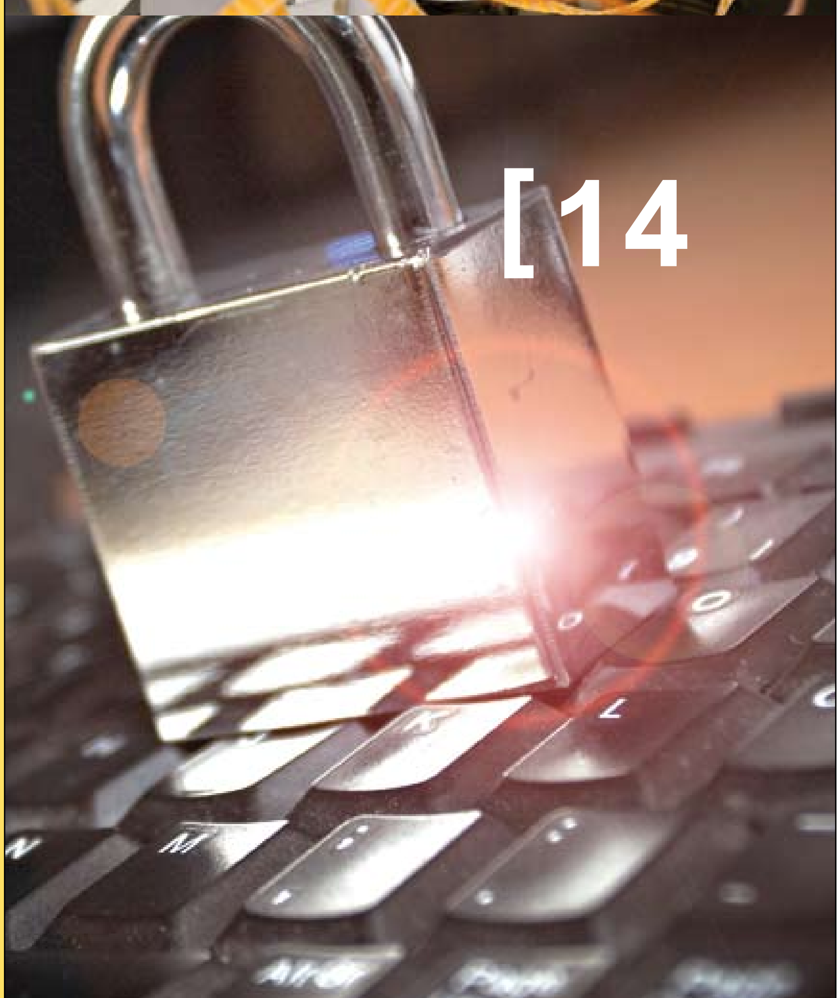

/contents

Jan van der Wel

CEO

/foreword

What is a solution worth if it does not solve the problem? As a developer of advanced products and systems, people ask us to come up with solutions. In order to develop a smart solution that works well, it is essential to know and understand how the problem has been defined and what the technical limiting conditions are. We therefore make discussing these issues with the customer a priority.

>>addicted to a solution

Unfortunately, we still regularly get project enquiries where a solution has already been chosen, even though the definition of the business problem behind the solution is not clear to us. That may be because there are company secrets that need to be kept, but time pressures or lack of clarity about the problem definition on the part of the customer may also be the cause.

Different interests or requirements generally lie behind the chosen solution. These may even conflict with each other or not be technically achievable. Solutions based on conflicting or unrealisable requirements can end up being too expensive or even cause projects to fail entirely. Being addicted to finding quick solutions to problems means solutions are often selected in haste. Those solutions then turn out not to actually be solutions, but instead create new and possibly greater problems.

So where necessary, we deliberately take some distance from the solution put forward and remain curious about the definition of the problem to be solved. For us and our customers, creating the optimum solution is an exciting voyage of discovery. By getting to grips with the question ‘what needs to be solved’, we are able to contribute our expertise and creativity and so optimise the success of projects for our customers. The article ‘problem identification comes before the solution’ on page 12 of this edition of Objective examines this issue in greater depth.

Happy reading!

/technology

GPU: brute processing power usable for more than just images

The GPU (Graphics Processing Unit) in a modern PC is a monster number cruncher. When producing a smooth flow of images, the GPU performs more calculations per second than the processor (CPU). But in practice, using GPUs for other applications is not so simple.

Driven by developments in the 3D gaming market, very fast graphics processors have been developed. In the first graphics applications, the PC's CPU calculated the image to be reproduced pixel by pixel. The image was then sent to the video card pixel by pixel, which in turn reproduced only the pixels. This is a highly processing-intensive task, certainly if this example needs to be performed for a 3D game at 30 frames per second with a resolution of 1920 x 1200.
In the nineties, ever more powerful video cards came onto the consumer market and the video card acquired its own processor, the GPU. Instead of directing the video card pixel by pixel, objects and settings are delivered allowing this separate processor to work independently alongside the CPU.

Standardisation

In 3D-gaming, a consumer market, very high volumes of these processors are sold. Because the producers of the GPUs and the games are different companies, the need arose for standardisation of the interfaces (APIs). The most important 3D interface standards are DirectX/Direct3D and OpenGL. What it comes down to in practice is that the GPU producers strictly follow the specified functionality of these 3D-interfaces (APIs). ATI (now AMD) and NVIDIA are well-known producers of GPUs.

The core of 3D games is the depiction of a virtual 3D world on the screen from the perspective of a (virtual) player, at a speed of for example 30 frames per second. This is achieved by performing a whole series of operations in a fixed order. This process is called a rendering pipeline.

History

As video cards became ever more powerful, an ever increasing proportion of the rendering pipeline was executed in hardware. The processing steps in the rendering pipeline are the so-called shaders. For example, a vertex shader adds particular 3D effects to objects, a geometry shader generates new objects based on previously defined objects and a pixel shader calculates the colour of a pixel. An important drawback was that by defining specific shaders in hardware, the calculating power of each shader was also fixed. Yet different games place different demands on the distribution of processing power between the different shaders.

Relatively recently, as part of the DirectX-standard, a Unified Shader Model has been specified. This is a processor model equipped with an instruction set that is suited to the functionality of all the different types of shaders. The advantage of this concept is that the full processing power can be deployed and is programmable at each step of the rendering pipeline. This new type of graphics processor also offers the first practical prospect of using the GPU for different processing tasks.
.

Processing power difficult to harness

The modern GPU has enormous processing power. It can handle up to 30 times as many floating point calculations as a regular CPU (see figure). That sounds great: so much more processing power for broadly the same price. Anyone would want that. But if something sounds too good to be true, it often is. There are plenty of ifs and buts attached to 'alternative' uses of a GPU, other than as a graphics card in a PC. The arithmetical problem needs to suit the architecture of a GPU, as does the floating point precision, and in practical terms it is not possible to combine a GPU with hardware architecture that is different to a PC's.

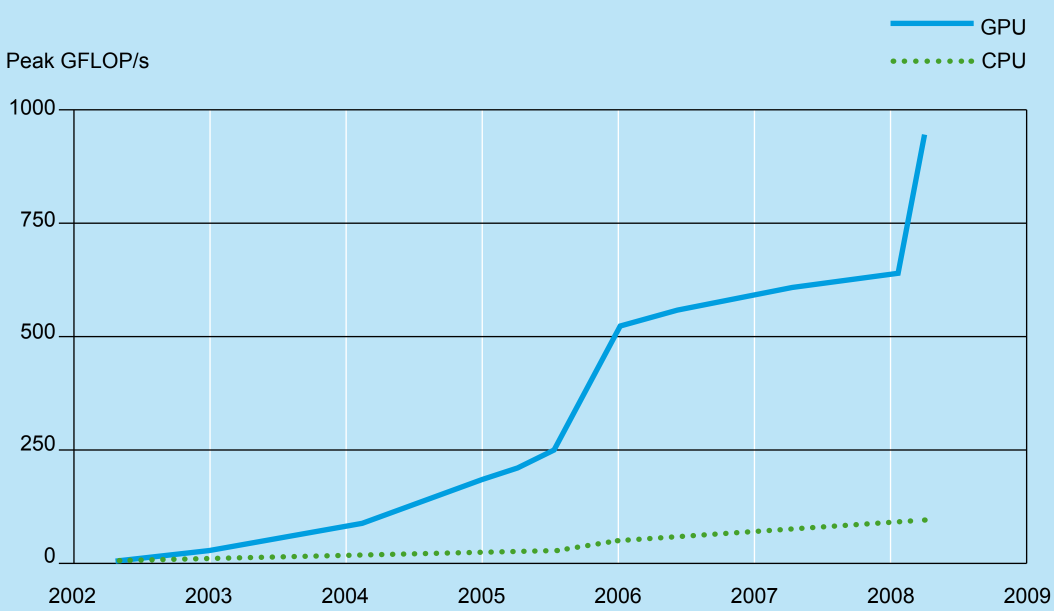

Dream scenario: GPU versus CPU
Optimum deployment of a GPU for image processing can yield the above curve (with many more MIPS for the GPU than the CPU). Beyond the specific graphics application, a GPU is limited. Running non-graphics applications on a GPU would show a different picture: the performance collapses to below the CPU-curve.

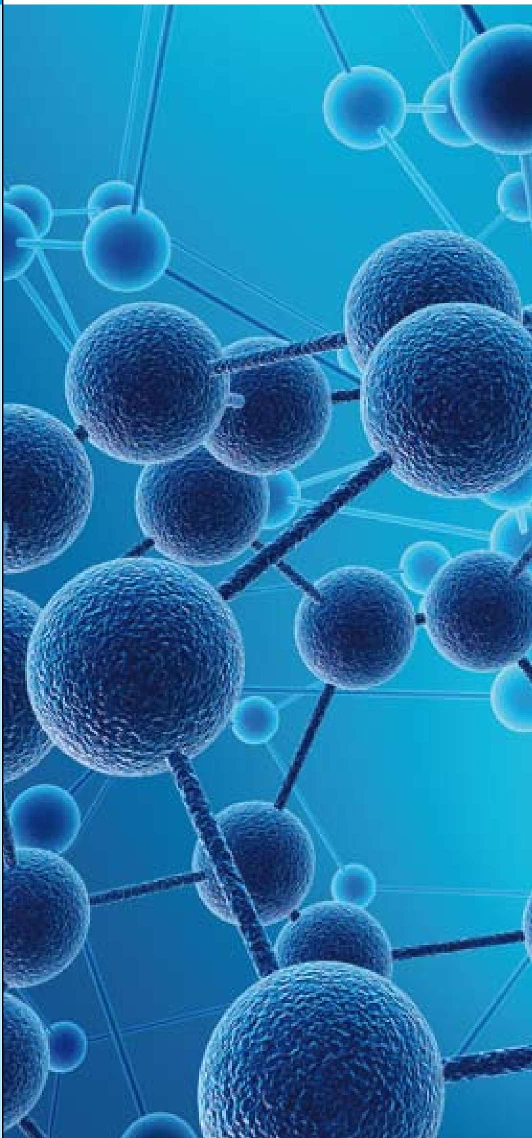

Deeper into the technology of the GPU

So should we simply conclude that while the images a GPU generates get better every year, it will never be more than a picture machine? Things are not quite so black and white. In certain cases and under strict conditions, a GPU could serve as an alternative for other big processing tasks beyond graphical work. To understand this, we need to delve a little deeper into the technology. A GPU is a parallel processor: it carries out a very large number of parallel calculations. This is reflected, for example, in the large memory bus widths, up to 256 bits. By way of comparison: the CPU relatively recently went from 32 to 64 bits. So a GPU handles a lot more data at a time. But the design is intended for one-way traffic: a chunk of data comes in with elementary image information; it makes some calculations for each pixel, then it pushes the results out the other end in the direction of the screen.

A GPU is bad at making choices, such as going through typical "if ..., then ..., else ..." software constructions. That slows it down. The point of a GPU is that it is designed to do the same for all parallel paths. This is called single instruction, multiple data (SIMD). There is an instruction decoder and a whole group of execution units. As soon as a choice comes along, it can only apply to one execution unit and the others cannot do anything effective. An ordinary processor is much better at solving if-then-else problems. It can smartly pick its way between a series of conditions and dependencies and save itself a lot of processing, like a skier continually making choices and steering to find the fastest way down a run. A GPU is more like a rowing boat with, say, eight rowers all trying to reach the same goal. If each rower were told to row a different course, the only way to achieve it would be to row eight different courses separately, each time with just one of the rowers rowing. So a GPU is good at parallel processing of arithmetical calculations according to fixed formulas.

Directing calculations using API

If a GPU is to be used, it will require special instructions, which can differ for every brand of GPU. The standard APIs are not usable. It is therefore essential that access to the GPUs is standardised, in order that the required functions can be provided at a reasonably abstract level. At the moment, the two most important APIs are CUDA from NVIDIA and the brand-independent OpenCL (Open Computing Language). OpenCL was initially defined by Apple Inc., who then turned it into an open standard via the Kronos Group. By using these APIs, the complex details of a GPU can be managed, allowing the programmer to focus on the problem at hand. Thanks to standardisation programmes (e.g. Matlab), libraries will also become available with GPU support for frequent problems, meaning that people will not need to reinvent the wheel each time.

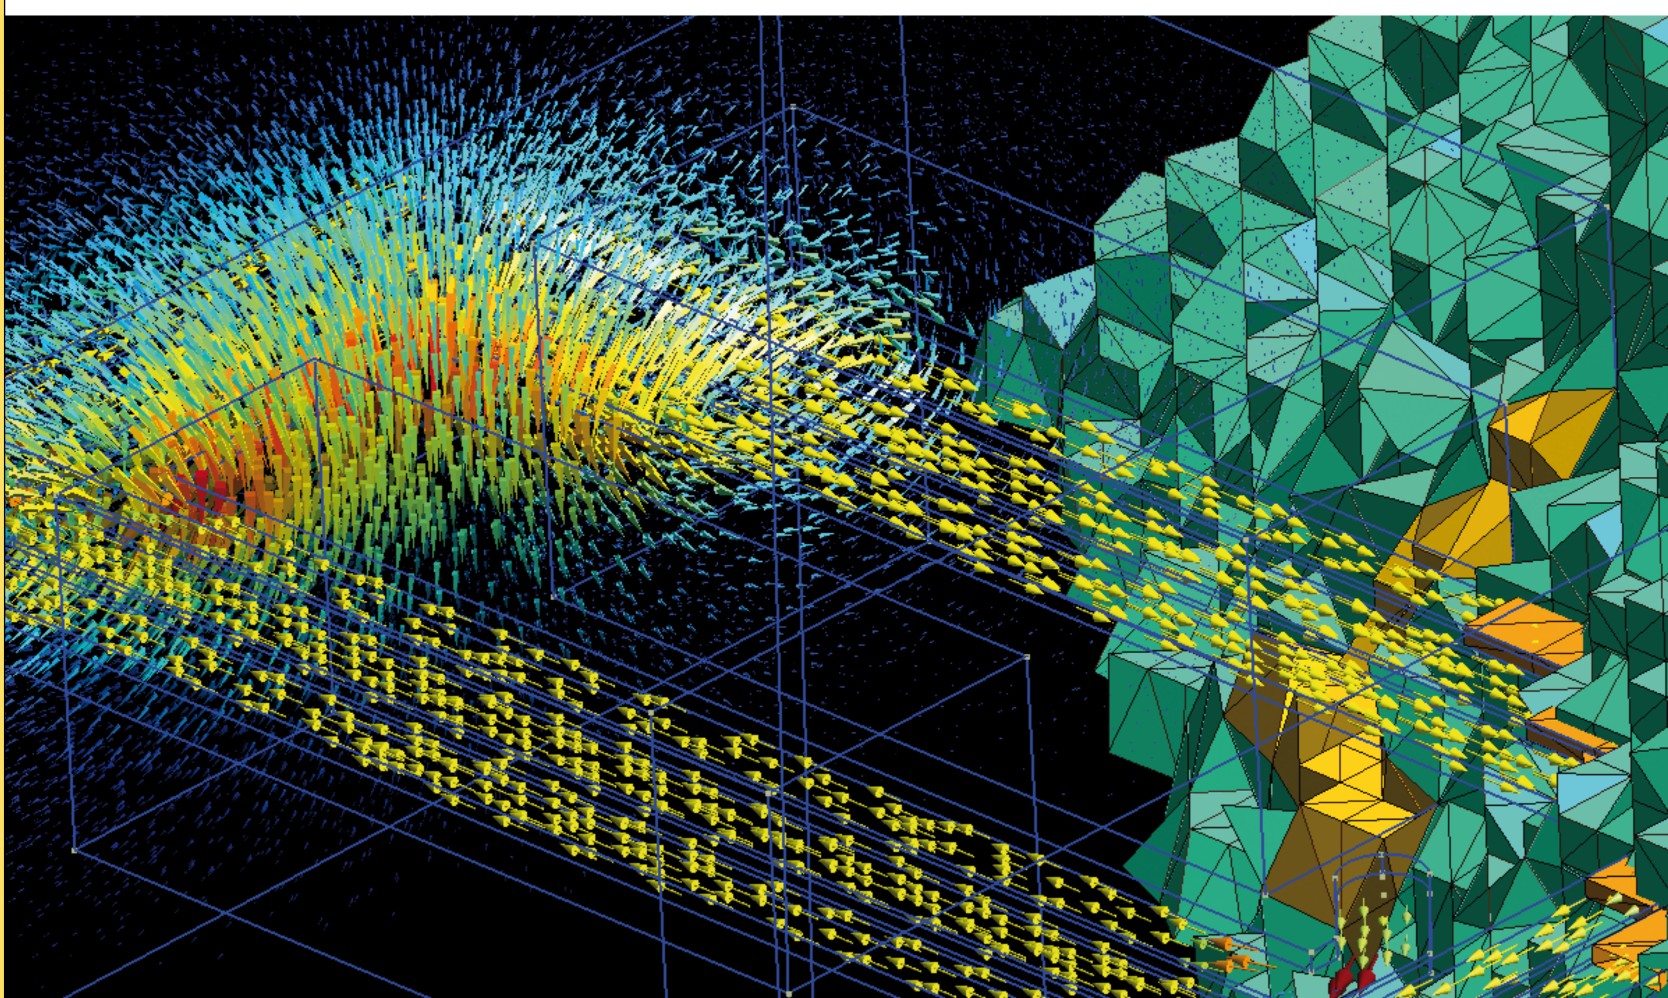

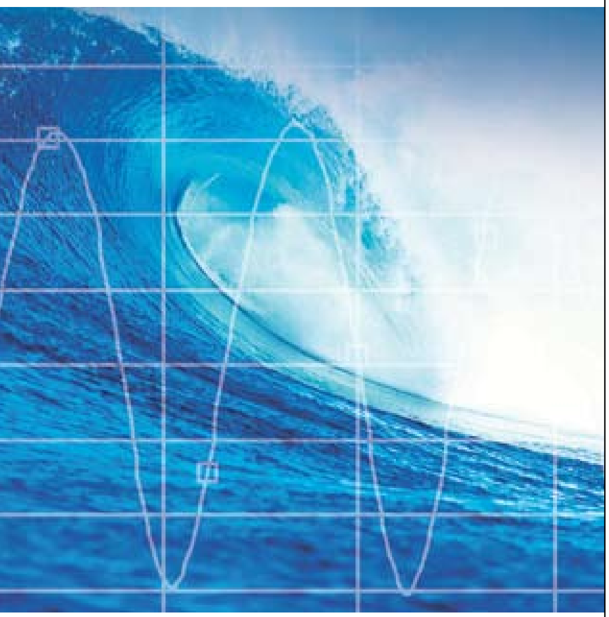

Using the processing power of the GPU differently

As indicated above, the APIs cited make it possible to successfully carry out suitable arithmetical problems, without output to a screen, faster on a GPU than on a CPU. Examples would include matrix calculations and Fourier transformations, for example in spatial problems such as 3D simulations of flows or electromagnetic fields. These are phenomena which may be analysed by dividing the world into small boxes or cubes to solve a physical equation (also known as the finite element method). Solving such equations requires a very large number of 'simple' calculating steps applied to a large data set, but also other large-scale calculations (particularly floating point) and image generation for the professional and scientific domains. For these kinds of applications, NVIDIA has developed a professional GPU known as Tesla. This allows a researcher to have a desktop which can compete with a supercomputer in terms of processing power for a relatively modest sum.

>>no "if..., then..., else..."

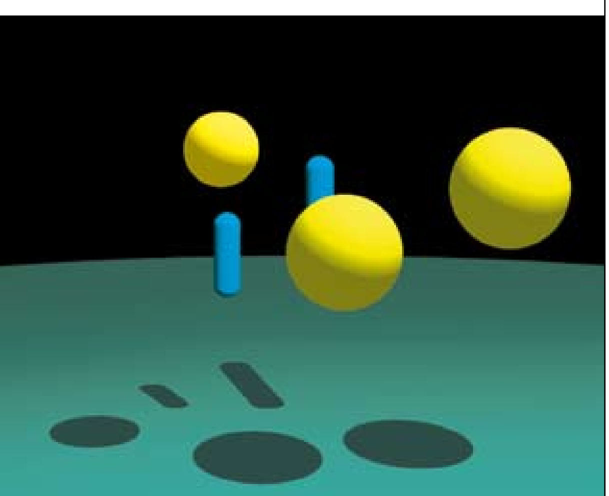

GPU in embedded systems?

For professional technical applications, a GPU is only usable in its natural habitat: on a graphics card, built into a PC, with standard software. Even where GPU chips are available separately, they are difficult to fit into embedded systems. The problem needs to fit with the architecture of the chip, the chip needs to fit into the system and the business model needs to be right. For instance, being designed for graphics applications GPUs have a short lifespan and as a result are difficult to source over the long term. The professional embedded world is virtually always better of with FPGAs, which will remain available for long enough. The choice of FPGAs is enormous: for every conceivable problem there is a specific FPGA somewhere. Not only that, they are highly flexible, because the developer himself can choose the balance between parallel and serial, between data load and processing power. In some cases not as many separate calculations are required, but they need to be performed on 10Gb data streams, such as in a high resolution camera. This generates a bulk of data that first needs to be processed before a PC can do anything with it. A GPU might be able to do the job, but for processing in the embedded domain Technolution always opts to process the data stream using an FPGA – precisely because of the issues discussed above.

The future of the GPU

One of the biggest limitations in parallel processing today is memory throughput speed. Larrabee is the code name of a development at Intel which marries an X86 CPU with a GPU to create a hybrid GPGPU (General Purpose GPU). In this case, that means genuine CPU cores with SIMD-support and a conventional cache. In the future, integrating a CPU and a Larrabee chip would allow both types of processors to be placed onto a single silicon wafer, giving very short and hence very fast connections. This would provide a lot of flexibility and enhanced performance for a broader range of processing-intensive applications. Intel's performance claims are being treated with scepticism by existing GPU specialists. But if this combination chip succeeds, it could herald the end of the separate graphics card and so bring entirely new applications for the GPU into view.

/customer interview

Optical lithography using electrons

A few years ago, the boundaries of optical lithography appeared to have been reached. Yet these days, the sky is the limit. Technically, much is possible, but costs are soaring. Which is precisely the area in which Mapper Lithography aims to compete.

Every year, the semiconductor industry develops chips with even smaller electronic switches. This is in keeping with Gordon Moore's Law, which states that technological progress will see the number of transistors on a computer chip doubling every 2 years. This requires production machines which can make ever smaller details. The core of the production machines is lithography: the transistors on the chip are created using optical techniques. Mapper Lithography is currently developing and building machines that write an image onto the silicon using not light but electrons, in much the same way that an electron beam scans the screen in an old-fashioned cathode ray tube. This technique was already experimentally used for lithography in the sixties. But the throughput time for electron lithography is low – very low. In the nineties, professor Pieter Kruit of TU Delft tried to ratchet up the speed by placing a large number of bundles in parallel. Two research students worked the idea out further. "That was Bert Jan Kampherbeek and myself," explains technical director Marco Wieland. "We graduated in the subject at the end of 1999. We decided not to take it to PhD level but instead to set up a company to really develop the technology properly. So in 2000 we founded Mapper Lithography."

Kruit's notion of using a large number of parallel bundles is still the heart of the current machine, as is his idea of switching the electron bundles on and off using light. The path of each electron bundle contains a switch that responds to light; so a light pulse switches the electron bundle on and off. Whereas optical lithography makes use of lenses, mirrors and prisms, the electron beams are steered using electron optics. "To make the optics, we use techniques from the IC industry", explains Wieland. "It's such precise work; it has gone beyond what you can do with mills and lathes."

Technical and economic limits

In the nineties, electron lithography was seen as the successor to optical lithography, which was running up against physical limits. However, those limits have been pushed back again and again with all kinds of tricks, but at the cost of making the machines ever more expensive. There is no longer a technical or fundamental limit to the development of optical lithography, but an economic one. And that is where the potential of Mapper's concept lies, believes Wieland. "Our machine costs five million and does ten wafers per hour. An EUV machine costs fifty million and does a hundred wafers per hour. In wafers per hour, that works out the same."

Because electron optics are much smaller than EUV optics, the heart of the machine is small: a cube of slightly more than a cubic metre. There will be ten of those cubes side by side on the factory floor. Including all peripheral equipment, the setup is a little bigger than the current generation of wafer scanners, but significantly smaller than the future EUV wafer scanners. So it will save on expensive clean room space. Mapper is well placed to compete with optical lithography in economic and technical terms. "But our big advantage is that we work without masks," emphasises Wieland. "In memory manufacturing, the issue is negligible; they make lots of the same memory chips with one set of masks. But in 'random logic' (everything except memory and standard PC processors), there are hundreds of millions to be saved using mask-free lithography. Ultimately, that kind of factory will be able to produce a lot cheaper using our machines." Due to the high accuracy but low throughput time, the machine will initially perform only the most critical mask steps, which is where mask costs are highest. Manufacturers will combine electron lithography with optical lithography (mix & match).

Photo above:
Technical director
Marco Wieland

Illustration:

1 Electron source

2 Collimator lens

3 Arperture array
Condensor lens
array

4 Beam blanker array

5 Beam stop array
Beam deflector array
Projection lens array

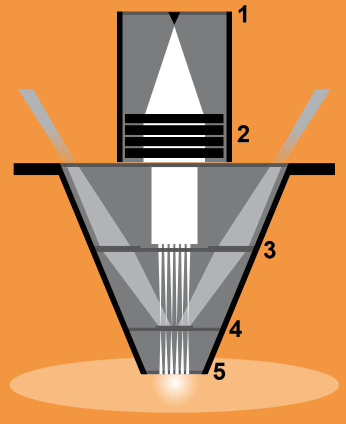

Industrialisation first

In 2007, Mapper reached an important milestone: proof of lithography. "We demonstrated that we could perform lithography using this principle. We did so with an experimental machine that contains the core of the Mapper: 110 electron bundles that we switch using light, electron optics and data transport. We then decided to put pure technological progress on hold for a while in favour of first building an industry-ready system that meets the standards of the semiconductor industry. Only then can we ship machines to customers. There we can scale up experiments and make use of wafer pre-treatment and post-treatment plus analysis and inspection equipment. Only when the machine has proved itself in a proper wafer factory will people believe that it can work. That is important in marketing terms too: we want to prove that we are more than a bunch of people with a cool experimental solution. No, this is a real machine."

13,000 bundles

"There are three important steps we need to take in order to get from the current test machine to the final version. Our current prototype can only light up separate small fields. By building in a metrology system, we can move the wafer under the bundles and knit the fields together. This so-called stitching will be ready early next year. We then increase the speed per bundle by a factor of fifty. That will make writing a whole wafer a realistic proposition (in approximately ten hours). And thirdly, we are going to increase the number of bundles to 13,000."

Cooperation

Electron optics is Mapper's core competency. For the other competencies, such as electronics and mechanics, the company works closely with partners. For example, Technolution supplies a significant share of the electronics. "The operation of the electronics (the data path) is closely linked to the machine. What we like about Technolution is that they are good at contributing ideas at the system level of the machine. That speeds things up significantly." Early on, the data path was the big bottleneck. Every second, fifty terabits of data go to the electron bundles. That is the equivalent of the data on a thousand DVDs, every second! And the transfer has to proceed with extremely high reliability: not a single bit must be missing, because that could cause part of the circuit on the chip to fail. By cleverly placing sufficient channels in parallel, that bottleneck was resolved.

Confidence

Although it will still be a few years before the first complete machine rolls out of the factory, confidence in the product is high. Future customers are already generating income. "They want the product to become a reality and they are prepared to invest in our research", explains Wieland. For the mass production of its machines, Mapper will seek collaboration with another machine manufacturer from the semiconductor industry. "It doesn't necessarily have to be a litho manufacturer", says Wieland. "We think the time will be right in 2010, once we have technically demonstrated the stitching and are able to achieve the required bundle flow and numbers of bundles in separate tests."

/case

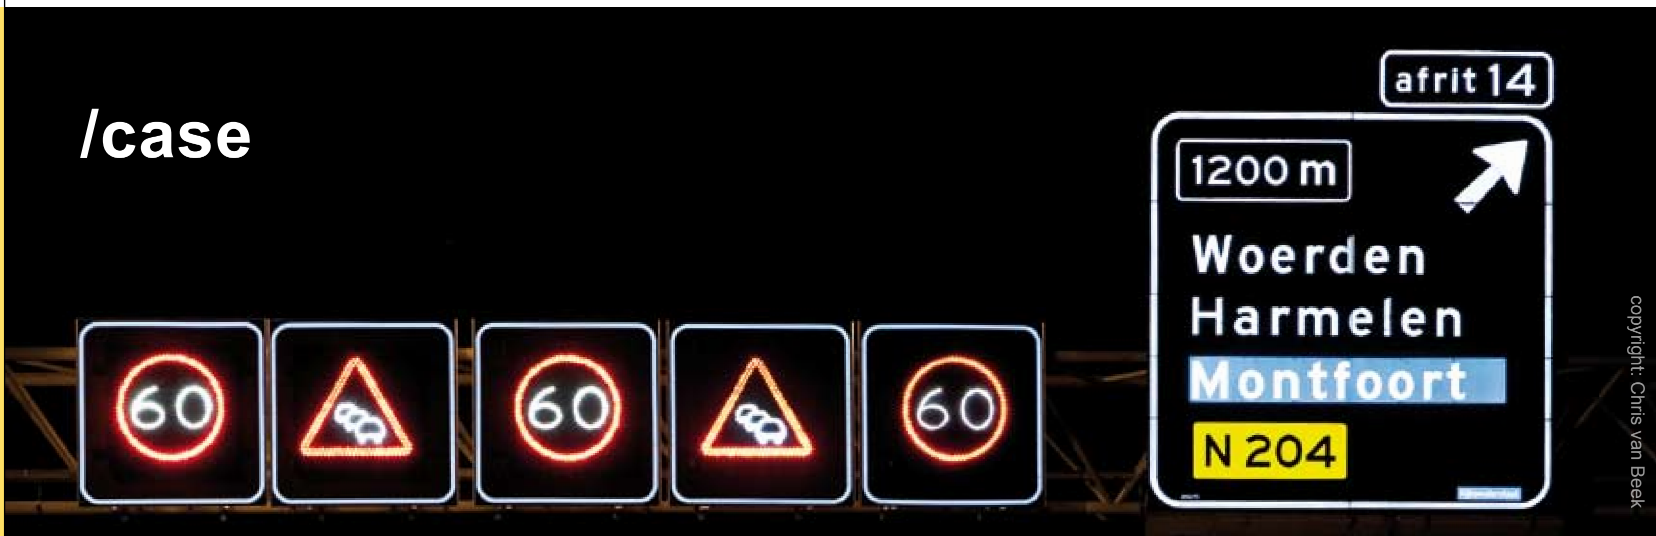

Drive faster when you can, slower when you have to

A speed limit that suits the conditions on the road, isn't that what every driver wants? Faster in the night-time when it is quiet and not quite so fast when it is busy or the weather is bad. The Dutch Directorate-General for Public Works and Water Management has started tests with dynamic speed limits on three different stretches of motorway in order to better understand "tailored" speed modifications. These tests are not only about winning support among drivers. The Dynamax project involves maximum speeds being tailored to actual traffic, road and local conditions at different locations. The tests are designed to reveal what the effects of this new way of handling maximum speeds will be on traffic flow, the environment and safety

Shorter journey time

On the A1 at Bussum, the Directorate-General is investigating what happens if the speed limit is raised to 120 kilometres per hour during quiet periods. The Dynamax system measures traffic intensity using vehicle loop detectors set into the road surface. If the volume of traffic allows, the roadside signs show the symbol for "end of all restrictions" and motorists are allowed to drive at 120 kph. When it gets busier, the signs switch back to 100 in a red circle.

Limiting environmental impact

The environment plays the decisive role in determining speed modifications on the A58 at Tilburg. Whenever the concentration of fine particulates is set to exceed pre-determined levels, drivers have to slow down. The input data is supplied by the Royal Netherlands Meteorological Institute, which calculates a five-day forecast based on a model. If the maximum value is expected to be breached, the Dynamax system reduces the maximum speed to 80 kilometres per hour and this is shown on the roadside signs. The National Police Agency is also informed, so that it can enforce this speed limit.

Improved road safety and better traffic flow

On the final test section on the A12 between Bodegraven and Woerden, two tests are running which are both aimed at improving road safety and enhancing traffic flow. Speed restrictions are introduced in case of rainfall – in fact before any rain falls. The Royal Netherlands Meteorological Institute's rainfall radar supplies the input data for the system. The speed restrictions remain in force for a time after a shower, because a wet road surface reduces visibility and increases braking distances, leading to a higher risk of accidents.

On the same section, a study is underway into the possibility of preventing so-called 'shock waves', a well-known phenomenon to traffic scientists. Car drivers are also familiar with the phenomenon of apparently spontaneous short traffic jams. Their cause is often a small disruption in traffic flow when the volume of traffic is increasing. For example, a driver braking hard can cause a ‘queue wave’. This ‘wave’ moves against the direction of travel and it can be a long time before it resolves itself naturally.

The algorithms in the Dynamax system process the speed of passing vehicles and detect such shock waves based on this data. As soon as a shock wave is observed, successive maximum speed adjustments are made over several kilometres in an attempt to dissipate it.

Challenging task

It takes a lot of work to make tests like these succeed. Technolution has been asked by the Directorate-General for Public Works and Water Management to develop the technology to manage the roadside infrastructure from traffic centres.

The Dynamax system lies at the heart of each of the tests because it converts input data into a plan of measures – a plan which determines what sign should display what speed limit at what time. In order to display speed limits, the Dynamax system makes use of the Motorway Traffic Management system (MTM) at the traffic centre and the roadside equipment linked to it.

Together, these systems are responsible for ensuring that road users are not shown confusing or erroneous combinations of images which might endanger road safety. For example, the Dynamax software makes allowances for the fact that the various algorithms may want to display signals to road users at the same time. It also allows the regular queue detection – which warns traffic of queues ahead – to continue to do its work independently of Dynamax

Technical implementation

The tests will provide plenty of challenges in both technical and organisational terms, not least because Minister Eurlings himself is the initiator of the Dynamax project. The way in which the technology is implemented and how the wishes of a large group of players involved are accommodated are at least as important as the technology. Those involved include the traffic experts at the Directorate-General for Public Works and Water Management, who define the basic conditions: which speeds need to be displayed under which conditions. Another interested party is the Directorate-General’s test centre in Delft, which tests whether the system meets all requirements, with a particular focus on road safety. In addition, there are naturally many points of contact with the customers whose job it is to implement the project on the roads and with the managers and suppliers of the systems at the traffic centre.

But because the tests are taking place in the real world, the regional traffic centres and the road traffic controllers who work there also have their own requirements. The automatic systems must seamlessly fit with the way these traffic controllers work – after all, they have their own responsibility in managing the traffic in their region on a day-to-day basis. In this context, the A12 trial is a particularly interesting case because the test section extends across the control area of two traffic centres: Zuid-Holland and Utrecht.

All in all, the Dynamax project proved to be a challenging commission and one which had to be carried out under significant time pressure because the start of the tests was already fixed at an early stage. An important factor in the success of the project was our extensive experience with dynamic traffic systems, both at the roadside and in traffic centres. As a result, we were able to quickly translate the commission into clear system specifications. This and a results-focused, no-nonsense approach were key to a timely and satisfactory completion of the various tests. Within six months of the start of the project, the Minister opened the test on the A1 and now, a little more than a year on, all the tests from the original projects are operational.

Inspection is in full swing

It is early days yet for results. All tests are accompanied with investigations. Rijkswaterstaat measure the speed and flow of traffic, and specialized agencies research the impact on behavior, safety and environment. It is certain the tests will be continued on the A12 near Voorburg and the A20 near Rotterdam, concerning improvements of traffic flow as well as improving air quality.

Problem identification comes before the solution

A good solution stands or falls with proper problem identification. People often start thinking in terms of concrete solutions even before their demands and needs have been properly defined.

People are often inclined to think in terms of solutions straight away. Someone complains that they are feeling out of sorts and you tell them to go out for a walk – whereas the solution might have been a listening ear. But that is difficult: listening and continuing to ask questions without already working towards a solution. The business world also displays such human traits. A customer will often already be thinking about what he wants in very technical and product-specific terms. But his demand is based on something, on a need. He wants to improve or speed up his processes. To achieve a solution, it is important to understand the problem behind the demand. This will make sense of the wishes and requirements expressed and help us develop a solution that meets that need. The essence is not to fix the "how" (for example, a system for image processing using 8 processors) but the "what" (a system that can suppress noise in incoming images by 10% at 60 frames per second with an image size of 4096 x 4096). This is called functional specification. And it starts with good problem identification that answers the "what" and "why" questions before it addresses the "how" and "when" questions. That answer will depend on the core competencies of an organisation. The core competency of the Dutch Directorate-General for Public Works and Water Management is water and transport. In the past, it would order an overhead sign gantry, specified down to colour and material. Nowadays it will ask for a solution to reduce the pressure of congestion on a road section. The thinking (and commissioning behaviour) has shifted from systems, through organisational processes, to functions or services.

Problem identification as a process

The package of requirements for a tender is more than a weighty document to which every party involved can contribute their own list of requirements. Coherence is important; instead of separate lists, what is called for is a single, optimum list. This requires consultation between all the disciplines involved: mechanics, physics and software, but also finances and logistics. The functional manager, also known as the system architect, stands between the product and the individual disciplines and seeks to establish the optimum package of requirements. Preferably in consultation with both parties, so that each can indicate what is and is not feasible or realistic based on their own expertise. This serves to create acceptance in all parts of both organisations: people feel heard and needs are met.

Culture
A customer must dare to delegate and overcome his fears – for example that his intellectual property (IP) will end up in the public domain. Because when IP is fenced off, it makes things more difficult. Creating an optimum platform to enable extremely fast processing of unfamiliar image processing algorithms is impossible if you do not know the functionality of those algorithms. It would be like having to make a frame for a painting whose dimensions you did not know. With a little openness, a much more suitable interface might be possible. A closed attitude puts a brake on creativity and makes insufficient use of the expertise available.

But a high degree of openness can lead to insecurity: “If we let you look that deeply into our organisation, why would we outsource the work? Why not do it ourselves? And what are our core competencies anyway? What do we want to be good at?” Existential questions, in other words, which can also arise in individual staff members. For example, the specialist who feels put on the spot when an outsider asks critical questions. Why didn’t he think of that himself? A good company culture strengthens the self confidence of staff. If they feel heard and supported, they will understand that the interests of the organisation and their own interests coincide.

Problem identification during development
Problem identification does not end after a package of requirements has been drawn up at system level. The complexity of systems and products is such that breaking them down into subsystems and components is essential. This involves solving problems at the level of subsystems or components. In order to be able to flexibly meet the wishes of the customer, an incremental development model is useful. This involves the developer repeatedly giving feedback to those concerned in order to make sure he is still on the right track. Short “feedback-loops” with prototypes, models or partial designs allow him to quickly gauge responses, exchange user experiences and create suitable solutions at the level of subsystems and components. He builds in flexibility by making the right choices in architecture, design and programming language at the outset. “By doing such and such, I am keeping these choices open for you.”

Verification at all levels is essential. For example, does the solution match the core problem as set out in the package of requirements? Is this solution essential, or just “nice to have”? A well structured development process and capable project management offer traceability: one can always see which needs have resulted in which choices. So it is possible to check that everything is still as it should be.

Conceptualisation
Even when all the wishes and requirements are neatly down on paper, things can still go wrong due to differences in conceptualisation. For example, a “user-friendly” system can mean all sorts of different things to different people. The point is to be specific. So functional specification is also about making sure everybody’s ideas match. Do we mean the same thing? For example, you might specify “user-friendly” as: no more than four choices per screen, in order to limit differences in interpretation. Or: you must be able to reach your destination within three clicks. Prototypes also help to ensure that visions match. Good problem identification is an essential start to any project. If the parties communicate openly on the basis of “what”, “how”, “why” and “when” and work together constructively, the most suitable and valuable solution will be the result. Everybody’s happy. The service provider gets to put its expertise and creativity to full use. And the customer? He gets a solution that matches his need.

'Everyware': always online everywhere

Thanks to smartphones and netbooks, the Internet has gone mobile. We can be online anywhere and at any time. Mobile Internet offers wonderful opportunities, but beware of the risks connected with privacy, security and reliability.

Being online at all times and places means being reachable, available and traceable anywhere. The GPS functionality built into a lot of equipment makes location-based services possible. For instance, Google can offer search results that are relevant to your location. Completely new applications are being developed with names like Layar, WikiTude and ShopSavvy that link technologies and databases together. Imagine, you are standing beside a beautiful canal in Amsterdam looking at a historic canal house and you wonder what it costs. You point your phone's camera at the building. The application in the phone links GPS, compass and camera image together so it knows where you are and which way you are facing. Next it accesses the Funda real estate database – unfortunately, the house is somewhat out of your budget, but you can still check out some photos of the interior. The tourist information database tells you that the church you can see off to your right is the Nieuwe Kerk, and that trams 5 and 9 for Central Station stop at the next bridge.

Outsiders break the rules: iPhone & Android

It is noticeable that many groundbreaking developments are coming from outsiders. With the launch of the iPhone in 2007, computer manufacturer Apple brought about a massive change in the telephony market and the Internet company Google may yet turn the market for mobile phone operating systems on its head with Android. This operating system is free for any phone manufacturer to use. Google earns nothing from Android, but the better people are connected with the Internet, the more advertisements Google can sell.

SpitsScoren

Although these developments are primarily driven by the consumer market, the business market can free ride on them too. Some technologies can be used directly, others may be cleverly combined. An example is the SpitsScoren project for the Verkeersonderneming (a body set up to promote traffic flow in the Rotterdam conurbation), which is designed to reduce the number of traffic queues on the A15 motorway. The project is focused specifically on car drivers. In order to reach and enthuse this target group, a community has been set up. This allows drivers keen to avoid rush hour traffic to find each other on the SpitsScoren website. There they can arrange to carpool using the "Pooll" application, entering their journeys in order to make matches with others. Every participant gets a smartphone to communicate how they plan to travel to work that day. In order to make sure that people keep to what they have agreed, the GPS functionality in the phones is used to record their movements. This project integrates existing systems: GPS, broadband mobile Internet and smart systems in the traffic centre and roadside cameras to make sure no-one is cheating the system.

Good privacy safeguards

The SpitsScoren project has good privacy safeguards built into it. Participants give their permission for their journeys to be tracked and recorded purely for the purposes of monitoring and advice. After use, the data is destroyed. Many companies actually depend on data from others. They will offer to store all kinds of things for you online, such as photos, your housekeeping book containing financial details or your documents. Many a supplier of free services scans such data for business purposes, for example in order to advertise in a targeted way. This is at odds with the requirements of privacy. Access to your data can be used for personal gain. But in Europe, and therefore also in the Netherlands, privacy legislation is fairly strict. Partly for this reason, the smart energy meter was scrapped, a system which could have helped encourage consumers to save energy. It remains a question of balancing convenience and privacy: what is the consumer getting in terms of convenience and what information is he prepared to provide in return?

Mobility services

The mobility project SpitsScoren is one of the test projects linked to the imminent introduction of road pricing. Soon, every car will have equipment installed that can calculate the cost of the journey. Service providers will also be able to use this equipment to offer other services, such as information on nearby filling stations or guiding you to a free parking space. In order to get this market up and running, a national subsidy scheme worth a hundred million euros has been set up. Ultimately, these services will need to be able to pay for themselves. If you have something good to offer, people will pay for it. Such as reliable and practical information or tools that make travelling easier and more pleasant. Such convenience is still lacking for many logistical systems (navigation and parking assistance in the city, for example).

Everything online everywhere

In the future, everything will be connected with each other. From the fridge that orders cheese and milk by itself to the car that brakes because someone else has hit the brakes a few cars ahead. The vision of intelligent devices that are integrated everywhere, are permanently connected with the Internet and can find all the relevant data there is also known as “pervasive computing” or “everyware”. We have not reached that point yet but some of the trends are already discernible. We are becoming ever more connected and we are storing more and more data in the Internet cloud. Where exactly the information is stored is not the point, as long as you can get at it with an Internet connection. Often the process is unconscious - it is simply convenient to be able to work on documents, presentations and software together in this way. This places high demands on reliability. If your PC crashes, you can quickly sort yourself out a new one. But if your Internet connection is down, you are at the mercy of your provider. Solutions to this problem are already emerging – such as an application that simultaneously stores your data offline and can synchronise it later. So you can carry on working even if your connection goes down.

Beware

The safety and reliability of the Internet are determined in part by the security of the computers connected to it. Most PCs run on Microsoft Windows, an advanced and complex product for a large group of inexperienced users. Banks demand “adequate PC security” for Internet banking. But what is adequate? We may reasonably expect every Internet user to have an up-to-date virus scanner and firewall, and to install update patches straight away. Even so, hacking into a computer is easier than breaking into a house. While a hacker drinks his coffee, his PC scans thousands of computers for weak spots – from Utrecht, or from Peking. For companies and governments alike, this is something to consider when they start offering sensitive services to customers and citizens over the Internet. In short, there are still some obstacles to overcome on the road to “everywhere” if we are to avoid taking a wrong turn to “noware”.

/employee interview

Paul van Koningsbruggen

*from vision to practice

Paul follows developments in intelligent transport systems (ITS) like a visionary. He spots trends and assesses needs within the traffic and transport sector in order to be able to take the strategic next step in time.

“After I gained my higher education diploma in civil engineering, I worked as a road traffic specialist for the Dutch Directorate-General for Public Works and Water Management and the Netherlands Organisation for Applied Scientific Research (TNO). I studied concepts related to traffic management and translated them into practice. At TNO, I made the switch from road traffic science to the technical side, because you need to be able to put your ideas into practice. I had been in touch with Technolution since the early nineties and in 2006 I joined the company. Here I can really make that step to practical solutions.

Assessing needs

As programme manager, I determine what the customer wants and what we are able to deliver. Those two things come together in a quotation and a project plan. But I find the phase prior to that even more fascinating: mapping long-term developments. By studying policy plans and spending a lot of time looking around in the outside world, I sense which way the market is moving. I see themes emerging. I zoom in on the points that everyone wants to tackle and the bottlenecks you encounter when you do. I ask myself what my customer needs and what Technolution needs to resolve them. This helps us identify the knowledge and expertise we need to be developing. In order to come up with new concepts, it can be very refreshing to take a look outside your own world, to feed your mind by taking excursions into other disciplines. For example, how do they go about solving problems in a hospital?

Jobs and roles

My job title is programme manager, which primarily involves managing business development and commercial processes. In addition, I do a limited amount of consultancy work, because I enjoy it. I work on European projects because I also want to experience that side at first hand, and in doing so, I gain expertise for consultancy jobs. It is my way of keeping in touch with the world outside and keeping abreast of what Technolution needs to be investing in.

For me, a commonsense, pragmatic mentality is important. I think about concepts at a fundamental level and then I discuss them with my colleagues, who immediately home in on the weak spots. I try to be visionary, they ask “what do we have available right now, which technology can we use?” Technolution’s sphere of operations runs from giving advice at policy level to developing practical systems. That makes Technolution unique. There are not many companies that cover the entire column in that way.

Internationalisation

I found the conscious step of entering the UK market particularly fascinating. We started from nothing by making presentations to companies and official bodies. We were building up brand recognition and at the same time getting to know the UK market. We are gradually gaining a foothold in the UK and the idea is that this will ultimately result in a Technolution office there.

I regard our corporate culture as being very down-to-earth, human and strongly focused on innovation. For a curious person, that makes this a great place to work. Here I can continue to look for innovations, as long as they are feasible in practice. “After all, innovation is only innovation once it makes it onto the street.”

/colophon

Objective is a publication of Technolution B.V. All previous editions of Objective are available for download from www.technolution.eu/objective

Distribution
Controlled circulation for connections of Technolution

Technolution B.V.
Zuidelijk Halfrond 1
P.O. Box 2013
2800 BD GOUDA
The Netherlands

T +31(0)182 59 40 00
E info@technolution.eu
I www.technolution.eu