Delft Outlook

2004.4

Research & Education at Delft University of Technology

The Internet: No quality garantee given

Robots with a Sense of Touch

Capturing the World in a Globe

Checking Mirrors with sub-nanometer accuracy

TU Delft

Delft University of Technology

CONTENTS

COVER: Fireworks at the seaside Town of Scheveningen (The Hague), Summer 2003

Delft outlook

Volume 21, № 4

Delft Outlook is published four times a year by Delft University of Technology
ISSN 0926-7212

Editor
Philip Broos (editor-in-chief)
Maarten van der Sanden

Contributing writers
Joost van Kasteren
Bennie Mols
Bruno van Wayenburg

Editorial office Delft Outlook
P.O. Box 5
NL 2600 AA Delft,
The Netherlands
Telephone +31 15 278 5860
Telefax +31 15 278 1855
E-mail: delftoutlook@tudelft.nl
www.delftoutlook.tudelft.nl

Design & lay-out
Tieme Dekker, Cok Francken
MultiMedia Services, TU Delft

Photography
Delft TU, unless mentioned otherwise

Translations
Marcus de Geus, Delft
Mrs. Drs. Margaret Clegg BA, The Hague

Printing
DeltaHage BV, The Hague

At last robots with a sense of touch

New applications of force feedback

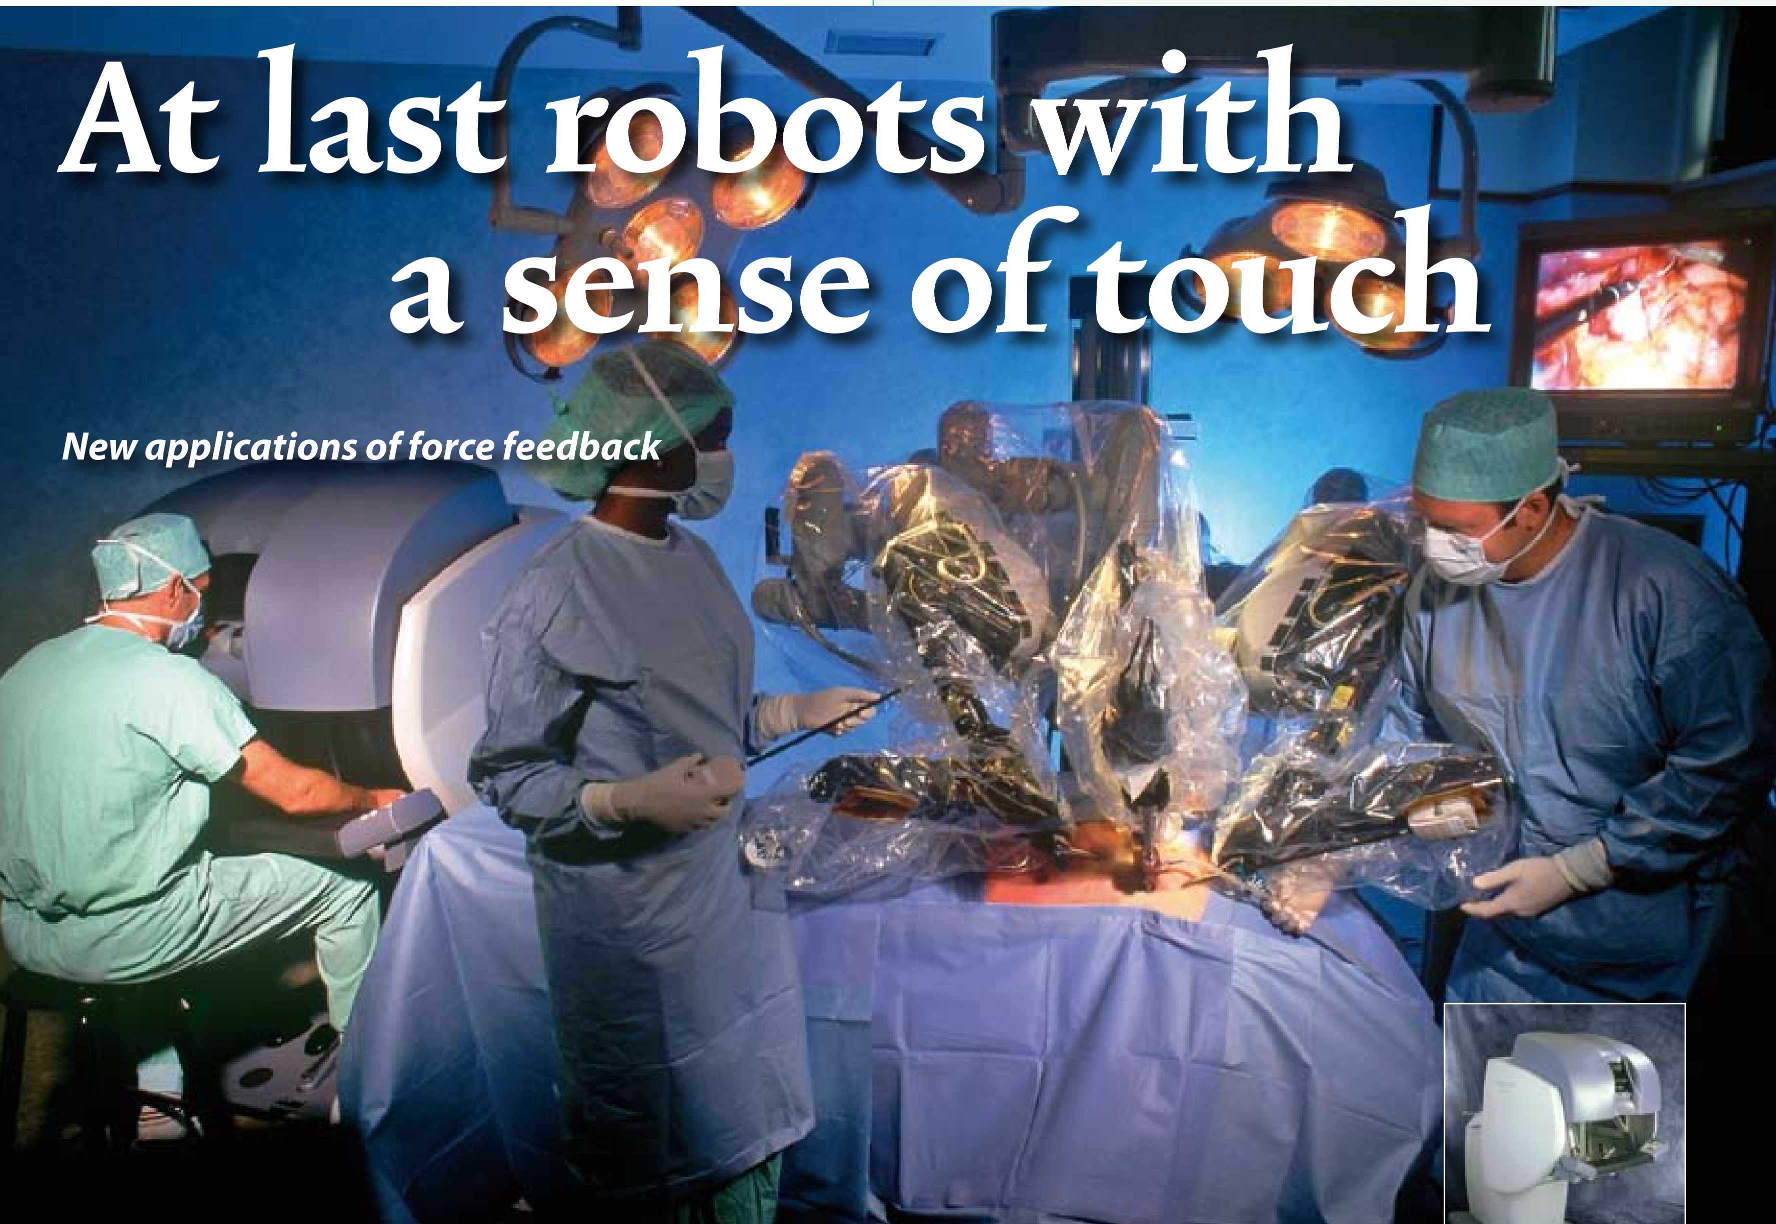

For a number of years surgeons have been using robotic arms to carry out precision operations. These technical appendages can operate without any tremors, and according to the manufacturers, they offer up to ten times the precision of the human hand. The downside is that surgeons using robot arms cannot feel what they are doing. Until now, that is. Countless research groups are looking into the possibilities of force feedback, i.e. transmitting the reaction force acting on the robot arm back to the surgeon. A lot of progress has already been made with soft structures, but hard structures remained a problem. Researchers at TU Delft have succeeded in building a first prototype that can grasp both hard and soft structures while enabling the remotely operating surgeon to really feel what he is doing.

By Bennie Mols

Seated at the console of Intuitive Surgical's Da Vinci telemanipulator, a surgeon controls the robot arms. On the console's display he is presented with a stereo image of the operating area inside the patient. The surgeon's hand movements are scaled down and reproduced exactly by the instruments held in the robot arms. However, the surgeon's hands do not feel anything of what he is doing in the operating area. The rest of the medical team can follow the surgeon's movements on the extra monitor (upper right). During the operation the telemanipulator is shrouded in plastic covers to maintain a sterile environment.

The control console of the Da Vinci telemanipulator includes two joystick-like grips which the surgeon uses to control the robot instruments. Feedback from the operating area to the surgeon is a purely visual affair.

Photograph: Intuitive Surgical Inc.

In 2001 surgeons at a London hospital operating on the prostate of a 61-year old man were the first to use three robot arms controlled by joysticks and a console in the operating theatre. The robot arms had to manoeuvre inside a one centimetre wide opening. Two small cameras attached to one of the arms provided a view of the operating area.

This type of operation requires maximum precision, which is something the robot arm is good at. The robot even filters out the natural tremors that affect even the best surgeons. The first robot operation in the Netherlands, on a gall bladder, took place at Utrecht University Hospital in late 2003. Elsewhere in the world, robot arms have even been used for basic heart surgery such as pericardial operations and heart valve repairs. For more complex heart operations human surgeons still outclass robots.

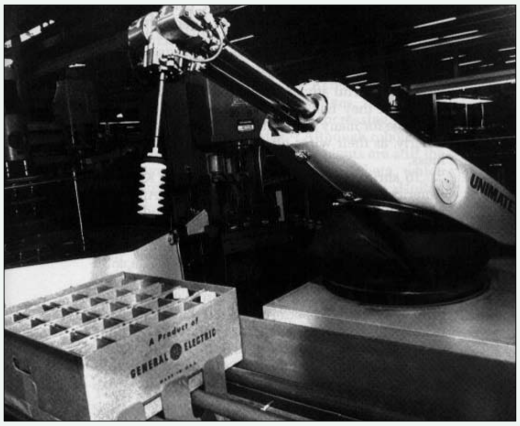

The first industrial robot, the General Motors Unimate, was introduced in 1961, and instantly reshaped industrial production processes. Today some 750,000 industrial robots are in operation, all of which are rigid, hard, and heavy.

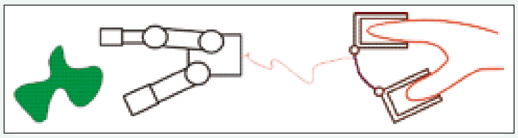

The classic remote surgery approach consists of a soft master interface (on the right) and a hard robot slave. This works well as long as the objects to be handled are reasonably soft.

An illustration of the contact instability that occurs when a rigid robot comes into contact with a rigid environment. The robot oscillates violently onto and across the contact surface. This behaviour is known as the hammer effect, and can even present a hazard to bystanders.

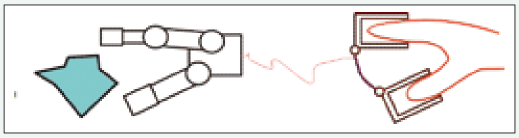

The classic remote surgery approach does not work with hard objects, because the forces acting between the rigid robot and the rigid environment build up too fast for the control system to keep up with. The result is contact instability.

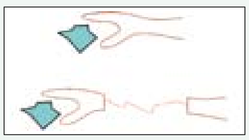

The ideal remote surgery system involves extending the hand of the operator with all its muscles and tendons. In this theoretical case the operator experiences the correct interaction with the object manipulated by the system.

The experimental set-up used at the Delft Biorobotics Laboratory to test the soft/hard theory. On the right is the master, with the slave to the left of it.

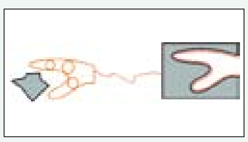

The approximation of the ideal remote surgery system as envisioned by the Delft Biorobotics Laboratory uses a soft slave hand that mimics the hand of the surgeon and a hard master hand that mimics the slave environment.

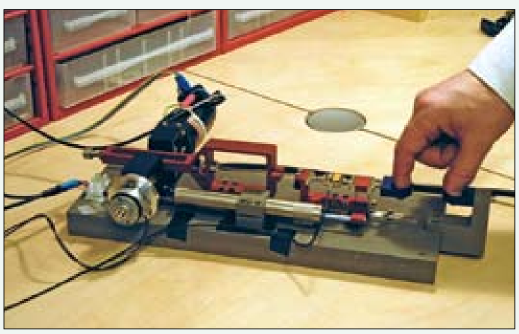

Close-up view of the master (the operator interface). The force is measured as close to the finger as possible. The system also measures the position and speed of the motor.

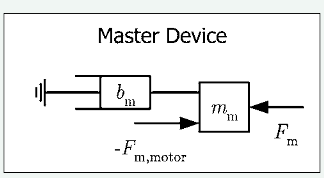

The master interface is approximated using a simple model in which the energy losses are represented by a viscous element (b_m), while the mass attached to the motor is represented by m_m, on which the forces of the motor (F_m) and the operator (F_m) act.

Movements in miniature

All remote surgery operations were performed using the Da Vinci machine, a commercially available robot, over two hundred of which have been sold. A surgeon operates it by holding a pair of grips that control the remotely operated robot arms which perform the actual operation. A video display provides visual feedback. Robot arms used in keyhole operations already have many advantages. They can remove surgeons' natural tremor, but they can also perform complex movements by exactly reproducing the surgeon's arm movements through a wrist joint in the instrument and the robot arm to which it is attached. A major drawback of the system is that the remotely operating surgeon cannot feel what is happening. He does not know whether he is operating on hard tissue or on soft tissue, since no contact sensation reaches his fingertips. To make the surgeon really aware of what is going on you need a force feedback system. When the instrument at the end of the robot arm comes into contact with a surface, the system feeds back the reaction force acting on the instrument to the joystick operated by the surgeon. This is also known as a haptic system. Researchers at the Delft Biorobotics Laboratory of the faculty of Mechanical Engineering and Marine Technology have been working on the development of just such a system.

The use of force feedback in robot arm control systems is not limited to medical applications. The European Space Agency (ESA) is developing robot arms for use on board spacecraft. Muscular dystrophy patients can be helped by fitting wheelchairs with force feedback robot arms. Force feedback robot arms are even used in robot rehabilitation. Force feedback enables us to control motor vehicles and aircraft using a joystick, which is a safe and relaxed way of operating things. Theme parks use the principle to create lifelike interactive attractions, and underwater operations would not be the same without force feedback systems. The TU Delft researchers are collaborating with a number of other parties, among them ESA, the Academic Medical Centre (AMC), and Exact Dynamics, Didam to ensure that their research follows real world needs as closely as possible,

Hammer effect

"In our laboratory we design robots inspired by biological functions," says Dr. Ir. Richard van der Linde, who supervises the research into haptic systems for remote control applications. Van der Linde previously developed the world's first autonomous two-legged robot driven by artificial muscles which needs less mechanical power than humans do. He is a part-time assistant professor at TU Delft, and is also employed by Altran Technology Consultants. In 2001 NWO, the Dutch Organisation for Scientific Research, awarded him the Veni, Vidi, Vici innovation impulse award for his research into biocompatible designs. The award money he received is now being used for the current research into force feedback in remote control systems.

Van der Linde: "The aim of biologically compatible remote control is to give the robot arm mechanical properties identical to those of the human hand. When you push your hand against something, it will flex while still enabling you to feel the difference between a hard and a soft surface. Even in contact with a hard environment, the hand remains stable."

This stability has proved to be a major problem when a robot arm is used to grasp a hard object while at the same time transmitting the sense of touch to the human hand operating the robot arm by remote control. Swedish doctorate student Göran Christiansson let me feel for myself what can happen if you use a traditional, rigid robot arm. I grab hold of the control arm and move the robot arm to pick up a hard object. As soon as the robot arm touches the object, it starts to hammer its surface like mad.

"Now you know why we call it the hammer effect," Christiansson says. "When the rigid robot arm comes into contact with a hard surface, the lightest movement results in a strong reaction force. The measuring and control systems simply cannot quantify the force and displacement involved: they cannot send commands back to the control arm in time to stop things going wrong. If the phase of the action becomes opposed to that of the measurement system results, the robot arm will start to jerk violently."

Crusher

Since traditional robot designers were aiming at accurate sensors, rigid structures, and fast control systems, instability upon contact with a hard surface became a common phenomenon. The lack of stability of the classic, rigid robot arm is not the only problem. If a rigid arm is used to handle a fragile instrument, it often ends up crushing the instrument when the robot lacks the

delicate sense of touch that humans have. As the robot arm makes contact with the object, the forces start to build up very rapidly because the robot is so heavy, rigid, and strong. This can also result in systems that form a potential danger to humans and the robots' surroundings. The contact between the human hand and the control arm is much slower because it is light, flexible, and soft. A system like that is much less capable of causing damage to its environment. Light and soft systems also tend to be less expensive to make.
The problems outlined above mean that the classic, rigid robot arm will only work well in contact with soft objects. In surgery for example, any contact with bone could spell disaster, while not all tissue is soft either. The surgeon may well be looking for a badly calloused section of tissue, or he may be using the robot arm to work with a hard instrument.

Dampened and soft robot arms

Bearing in mind the biological principles of the human hand, Christiansson, supervisor van der Linde, and supporting engineer Erik Fritz designed and built an experimental model that solves both problems. To begin with, both the operating end (the master) and the robot pincer end (the slave) have only one degree of freedom. When an imaginary surgeon on the master side uses his thumb and index finger to move two grip rings towards each other, an electrical signal transfers the movement to the robot arm, which has also been fitted with a pair of gripper rings that it can move to and fro along one axis. The single degree of freedom involved is sufficient to investigate the force feedback principle.
"When I take hold of a hard object between my thumb and index finger," Christiansson explains, "my fingers will flex slightly, but I can still feel that the object is hard. We are looking at ways of incorporating this underlying biological principle into the robot arm. That is the main innovation."
In technical terms the researchers opted for a pair of leaf springs with adjustable stiffness, with a variable damper set parallel to them. This is the first time that anyone has constructed a robot arm with variable rigidity and damping.
"We recently presented our concept at a conference," van der Linde recalls, "and some of the scientists there just couldn't believe that we were intentionally taking away the rigidity of a robot arm. The thing is that making the robot arm less rigid actually increases its application potential. It will enable surgeons to operate on both hard and soft tissue. Our solution is a full analogue of biomechanical muscle models."

Radioactive materials

The damper consists of a piston in an air-filled cylinder with a small hole in it, the diameter of which can be varied to adjust the degree of damping. The robot arm consists of a mass connected in series to a parallel damper/spring system, which in turn is connected to another mass. The result is that the robot arm can be made non-rigid.
Christiansson: "The question is, can you use a non-rigid robot to feel that a hard object is just that? You are inclined to think that this is not the case, but as it turns out, it works very well."
The operating end looks identical to the robot end, i.e. it has the same adjustable mass-spring/damper-mass system.
"Even so, the system is easiest to operate with the operating side made as rigid as possible," Christiansson continues, "producing a system that consists of a rigid operator arm and a non-rigid robot arm."
The asymmetrical set-up is a new and promising aspect of the development of robot arms for remote surgery applications. The aim is to make the robot arm a sensitive extension of the human arm. Previous development work on robot arms with force feedback to the operator panel does exist, in particular for handling radioactive material in the nuclear industry, but none included this asymmetrical way of working, using a non-rigid robot arm linked to a rigid operator arm. Instead of recreating the properties of muscles, traditional robotics used powerful but far from subtle electric motors.

Integrating mechanics and controls

When the slave unit developed at TU Delft makes contact with an object, sensors automatically measure the position of the end point and the force acting on it. The control system at the robot end then processes the position and force information, and passes it to the operating end in the form of an electronic signal. To supply the necessary power, both the master unit and the slave unit have been fitted with DC motors. So a motor supplies the power to push back against the human hand

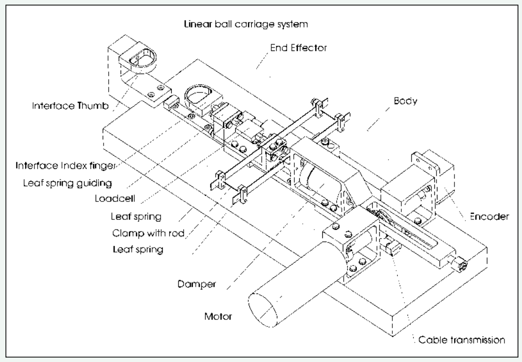

The slave-set-up.

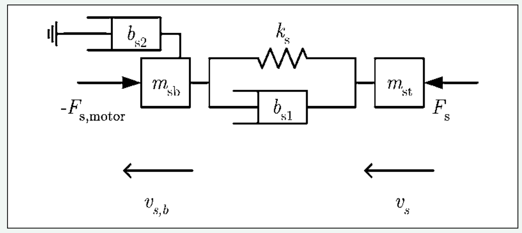

The slave interface is also being modelled using a simple model, with the energy losses represented by a viscous element (b_{s2}), and the mass attached to the motor represented by m_{sb}. In contrast to the master this system includes a spring (k_s) and a damper (b_{s1}) between the tip and the motor. These make the slave soft. Acting on this system are the forces of the motor (F_s) and the operator (F_s).

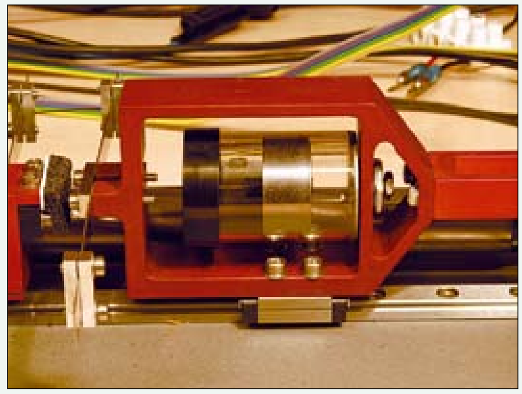

The spring (k_s) and the damper (b_{s1}) from the previous figure have been realised using an adjustable leaf spring structure and an adjustable piston and cylinder combination.

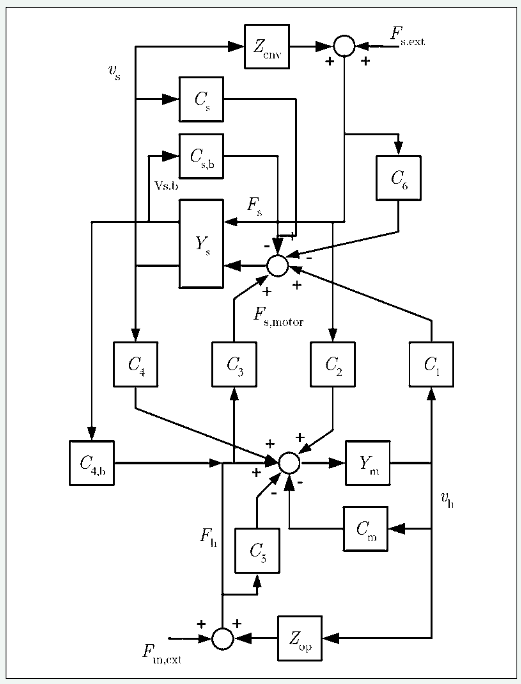

Although force feedback appears to be rather obvious, it requires a complex control system. Even a simple set-up with a single degree of freedom requires some 10 different parameters, as well as models of the environment, the operator, and the set-up. Since none of these models can ever be exact, finding the optimum settings becomes a complex matter which in practice often boils down to trial and error.

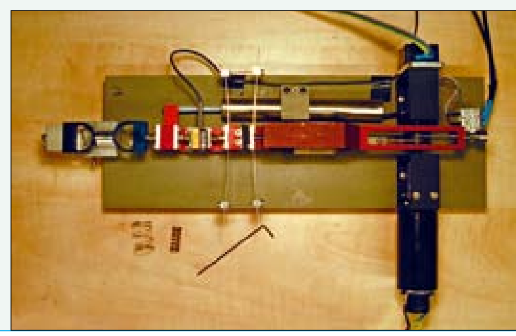

The soft slave developed by van der Linde's group features an internal element that gets deformed when a hard object is grasped between the gripper loops. This results in a natural build up of forces between the gripper loops of the slave and its environment.

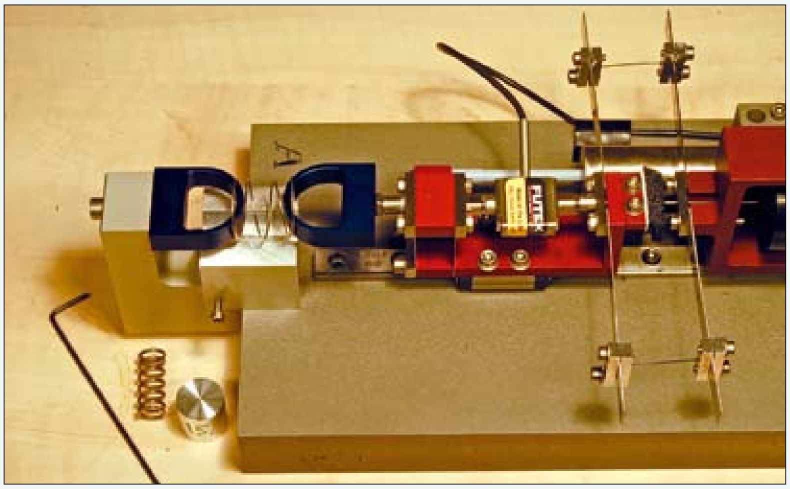

The soft slave can also handle soft objects. In this case, both the slave and the object become deformed as a result of the contact forces.

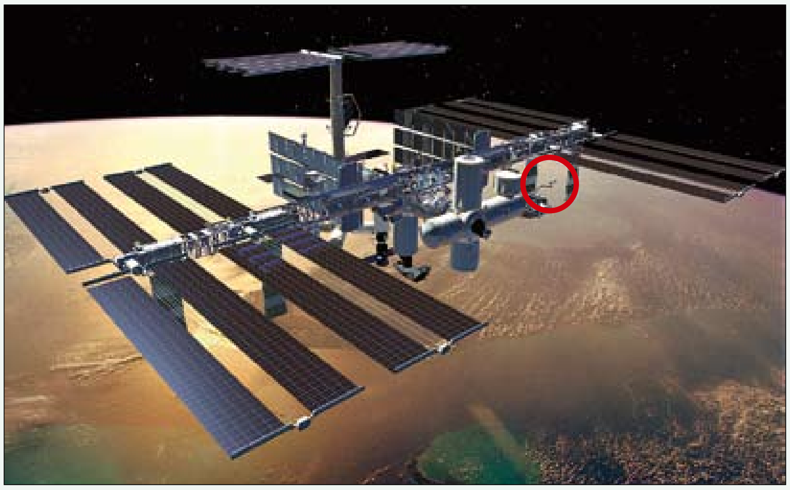

Illustration: NASA

The exterior of the International Space Station (ISS) requires regular maintenance. Since the robots developed so far (such as the European ERA developed by Dutch Space (marked by the red circle) are too unwieldy for maintenance work, astronauts have to venture outside for maintenance work, known as Extra Vehicular Activities.

The Eurobot is controlled by an astronaut from inside the ISS. Force feedback would be a useful feature to prevent the Eurobot damaging the ISS. For this purpose, TU Delft doctorate student André Schiele is working at ESTEC to develop an exoskeleton that will enable astronauts to operate the Eurobot while providing force feedback from the robot to the astronaut.

enabling the operator to actually feel the reaction force. The control system drives both motors. Ideally a human operator should feel exactly what the robot arm feels, however, the robot will operate better at certain frequencies. Christiansson: "At frequencies up to a couple of hertz we can achieve a force transfer of about 95%, which is very good. At higher frequencies the performance drops rapidly, but this is not a problem for most applications."

Two worlds

"In our system we have integrated the mechanics with the control system," van der Linde continues. "The Da Vinci robots now in use in hospitals try to solve everything using pure control systems, i.e. with electronic positioning and correction. My guess is that eighty percent of the world would like to solve robot control along these lines. Control systems are easy to use across distances, and to scale down large operator side movements to much smaller movements at the robot arm end. Even so, apart from the various mechanical alternatives there is still no force-controlled remote surgery system that is robust and fully stable in contact with hard surfaces. There is always the danger of contact instability, which is why we have added a mechanical damper and spring system to an electronic control system, combining the best of both worlds. Mechanics alone, just like pure electronic control systems, just cannot do this."

At the AMC in Amsterdam, fully mechanical master-slave manipulators are being evaluated.

"Surgeons tend to like the idea of a fully mechanical system," says supporting engineer Fritz. "An electronic control system uses motors that can add energy to the system. Most surgeons think that a motor has the potential to mess up an operation. With a fully mechanical solution the surgeon feels completely in control. So, we will have to win over the surgeons first before we can introduce electronic control systems in a remote surgery system."

Moreover mechanical solutions are still less expensive to manufacture.

Sound as additional information

Whatever the integration of mechanics and control system looks like, the information arriving at the operator end is always slightly filtered.

"One of the things we still know very little about," says Christiansson, "is which kind of information is most important for certain tasks. For example, what is the information you need to feel differences in human tissue? We don't know yet, but we will have to find out in order to set up our robot arm correctly. How rigid should it be, and how much will it have to be dampened? How sensitive must the sensors and motors be? How will time delays affect the force feedback? All of these are still practically unknown factors."

Christiansson, who took his degree at Chalmers University in Stockholm, is now halfway through his doctorate research. He intends to use the remaining time to find the answers to these questions.

"We will be asking various test subjects to perform experiments with our master-slave setup, so they can tell us what the relevant types of information are."

Fritz adds that if a surgeon looks at a display to see what he is doing he could at the same time be hearing what is happening to provide him with additional information that can help him assess the condition of certain types of tissue.

"It would enable him to integrate the information about what he can feel with what he is seeing and hearing. As a possible result, it could mean that the extra audiovisual information reduces the level of tactile information required," Fritz explains.

Flexible plastic fingers

Once the scientists understand what kind of information is needed to build the right kind of operator arm, the next step will be to create more realistic grips. Whereas current models still use a pair of loops through which you put a thumb and forefinger, future models may well use lightweight plastic fingers.

"At some point in the future they will become small medical instruments, flexible systems, that can be easily attached to a robot and so cheap that they can be disposed of after use," says van der Linde. The system will also have to evolve from its current configuration with a single degree of freedom into a mobile instrument with several degrees of freedom.

Van der Linde continues: "The crucial question really is, how good should good be? Even though you cannot see the forces, you can feel them. This makes it difficult to quantify what the human operator stands to gain from using robotic manipulators. Anyhow, the main point is that we have now managed to demonstrate using a highly simplified system that our concept with a non-rigid manipulator and a rigid control manipulator offers great advantages over the classic approach.”

According to van der Linde, remote surgery is a newly developing field, even though it has a history that goes back fifty years. And it can look forward to a great future both in medical applications, in the form of touch-sensitive instruments, and in other fields, including micro-assembly, the handling of radioactive material, and underwater robotics.

For more information please contact Dr. Ir. Richard van der Linde, phone +31 (0)15 278 6585, e-mail r.q.vanderlinde@wbmt.tudelft.nl, Ir. Göran Christiansson, phone +31 (0)15 278 5633, e-mail g.a.v.christiansson@wbmt.tudelft.nl, Ir. Erik Fritz, phone +31 (0)15 278 5633, e-mail e.c.fritz@wbmt.tudelft.nl, or André Schiele Dipl. Ing., phone +31 (0)71 565 37 60, e-mail andre.schiele@esa.int.

Additional information:
http://dbl.tudelft.nl
www.intuitivesurgical.com/products/da—vinci.html
ttp://robotics.estec.esa.int/
www.exactdynamics.nl
www.fcs-cs.com/robotics

Photograph: Exact Dynamics

Remotely operated mechanisms are also used for rehabilitation purposes. This picture shows a disabled person suffering from motor deficiency using the Exact Dynamics ARM, a mobile manipulator attached to the wheelchair. Force feedback might be very useful to this user group. In addition, the soft slave principle could improve safety.

Photograph: Exact Dynamics

People use force feedback in many everyday tasks. To perform the same tasks using a robot without force feedback would require lots of training.

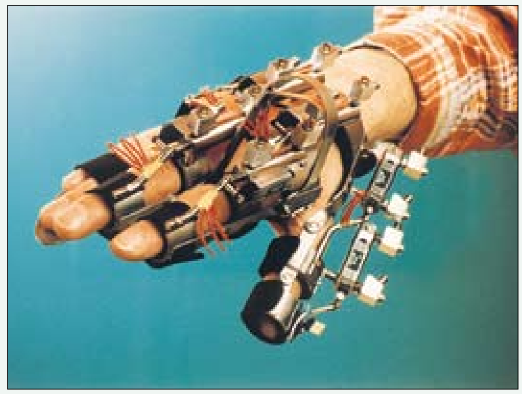

The human hand has 22 degrees of freedom and contains over 17,000 sensors. To reproduce the complete range of movements available to the human hand would require complex mechanics, as demonstrated by the hand from the European Teleman project from the 1990s (for which TU Delft acted as secretary; see also Delft Outlook 97.2).

Slower and softer, but still better?

The human hand cannot match a robot hand in many ways. Consciously, humans can vibrate their hands at no more than about five hertz, and even unconsciously, as when suffering from Parkinson’s disease, the human hand vibrates at no more than fifteen hertz. The precision of the sensors and actuators in the human hand is low, enabling us to recognise differences between one and ten percent at best. Doctorate student Göran Christiansson has me do an experiment.

“I will give you a little block to hold between your thumb and index finger. I want you to estimate its length. Then I will give you another block, and you will do the same. I then want you to tell me which is the longer one of the two. You may not look, just feel.”

I take first one, then the second block between my thumb and index finger.

“They feel the same to me,” I say.

“Wrong, one is three millimetres longer than the other one,” Christiansson laughs. “A robot would do much better.”

Humans can only feel forces acting with a frequency of less than 300 hertz. The neural delay between the moment the hand registers the force and the moment we respond to a stimulus generally is between ten and one hundred milliseconds long. From the finger, the stimulus has to travel to the spinal cord and back, and sometimes even all the way to the brain and back. The knee jerk reflex for example takes about one hundred and forty milliseconds. When a person accidentally touches a hot surface it takes up to half a second before they realize what is happening and pull back their hand.

“For a robot arm this kind of performance level would be abysmal,” says Christiansson. “It can easily outperform us on all of these counts.”

Nonetheless, the human hand is very stable and has a very finely tuned muscular system. A human hand contains some seventeen thousand sensors, although there is quite a bit of overlap. We use seventy percent of all these sensors to gather tactile information, and the rest is used to perceive pain stimuli.

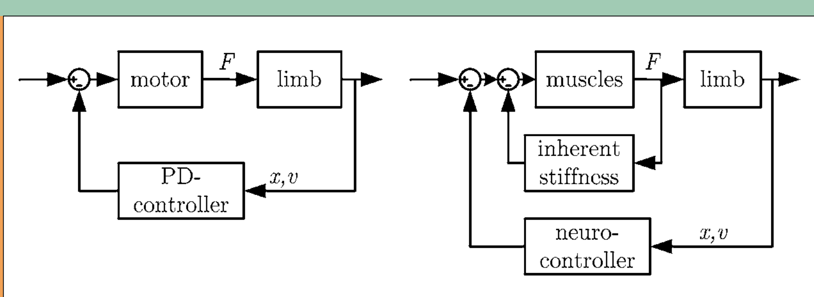

Robots — From big, heavy and cumbersome to small, light, and flexible.

The year is 1921 when Czech playwright Karel Capek introduces the world to the word ‘robot’. In his play ‘Rossum’s Universal Robots’ he describes how robots handle all the chores so humans can spend their time on leisure. It was 1961 before the first robot application appeared on an industrial production line when the U.S. company General Motors used a robot arm, the now legendary Unimate. The idea that humans could rest on their laurels and let robots do all the work for them has of course been relegated to the realm of dreams. Even so, the fact remains that robots can perform tasks that humans simply are not very good at, such as lifting heavy objects, or endlessly repeating the same actions, all with ultra-high precision. For decades traditional robotics focused on strong robots, producing large, heavy, rigid structures. The result was exceptional performance of power, speed, and accuracy, quite a contrast with the sensitive coordination and movements of the human arm with its damping and flexibility built into the muscles and tendons. Whereas our limbs can move lightly and elegantly, classic robot arms appear clumsy and move in sudden jerks. Also, the human nervous system includes a number of ingenious adaptive feedback mechanisms. The control mechanism of a traditional robot arm is essentially different from that of a human arm (see figure below). Biorobotics researchers are looking for ways of improving robot systems for specific tasks by incorporating biomechanical principles in the design of robot arms. The aim is to improve the system’s stability and efficiency and to build lighter and more flexible robots. In addition the biorobot arm is designed to give its operator a sense of actually being in contact with the object or surface the robot touches. This is one of the fields being researched at the Delft Biorobotics Laboratory.

BY BRUNO VAN WAYENBURG

Delft algorithm to save Internet from traffic infarct

Internet pathfinder SAMCRA for irritation-free Internet phoning

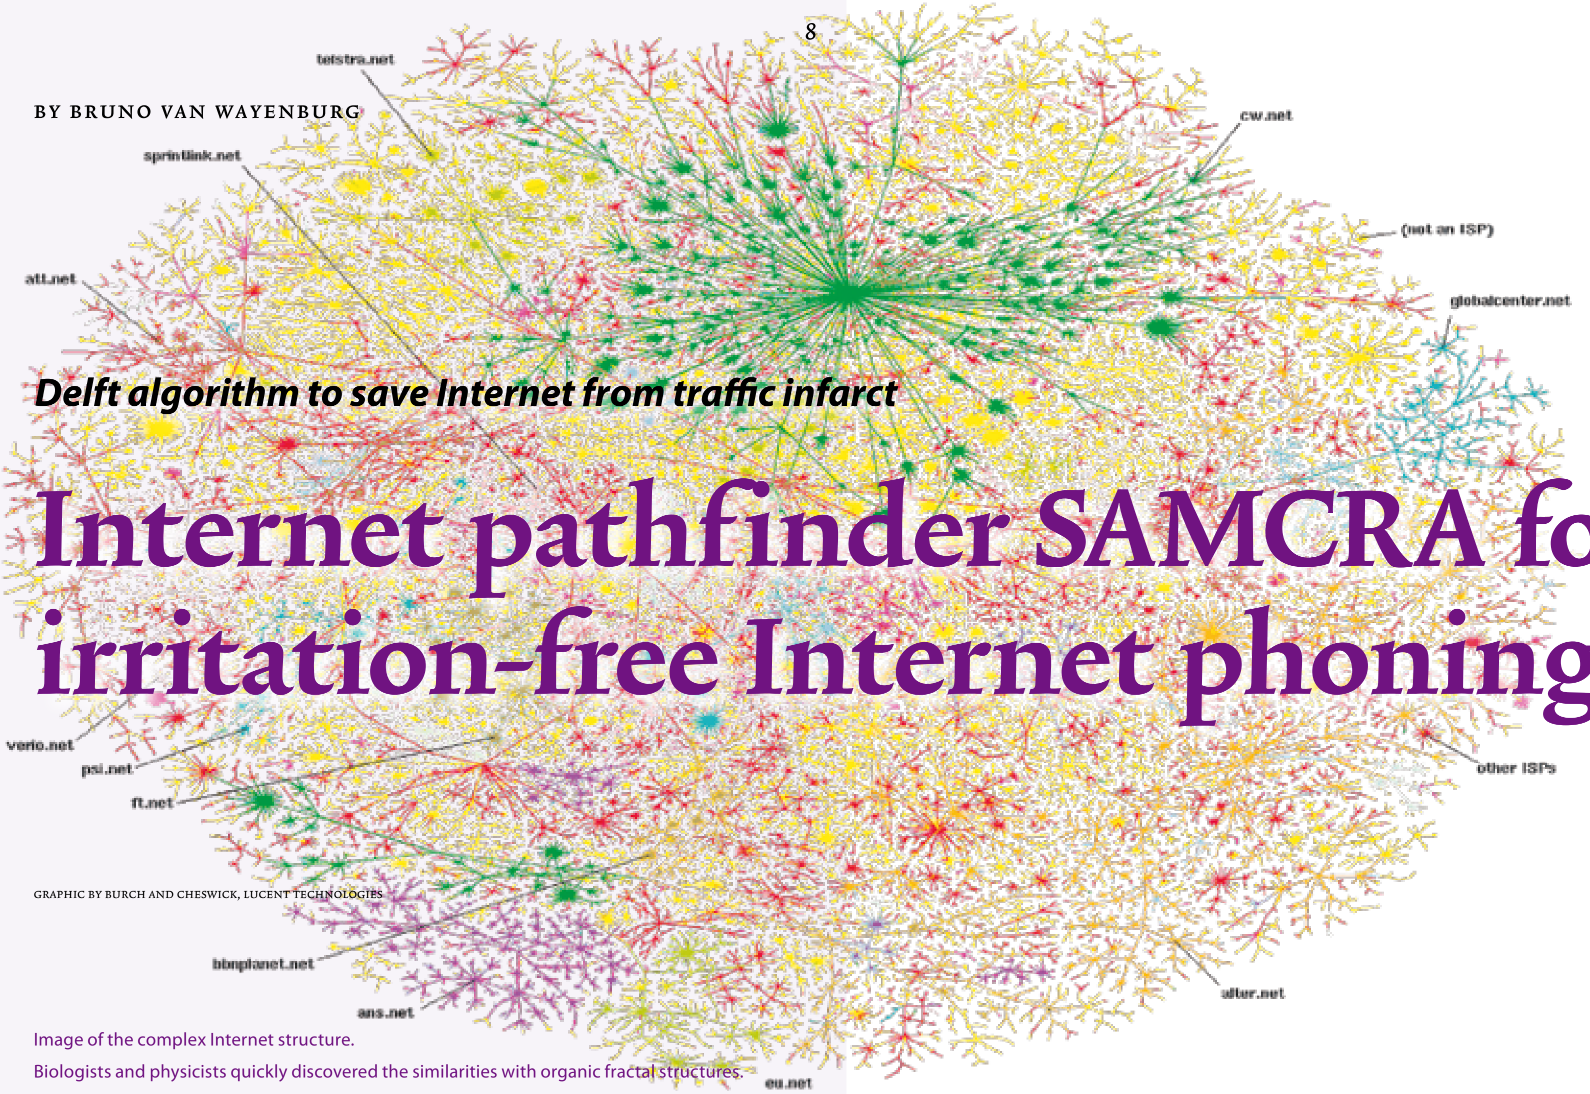

GRAPHIC BY BURCH AND CHESWICK, LUCENT TECHNOLOGIES

Image of the complex Internet structure.
Biologists and physicists quickly discovered the similarities with organic fractal structures.

The growth of the Internet is going to grind to a halt unless fundamental changes are made in its organisation, predicts Piet Van Mieghem, Professor of Telecommunication Networks at TU Delft. Routing, the information highway's traffic control system, is already a cause for concern. And, as in the case of road congestion, simply creating additional capacity will not provide relief. Van Mieghem's solution: simply offer better Internet connections for sale in order to guarantee the quality of the connection for Internet phoning or video conferencing and ensure the service proceeds without interruption. Internet service providers could then really start making money from the Internet. But this calls first of all for smart routing software which is guaranteed to find the best route through the network, preferably without having to compute for days. With the routing algorithm SAMCRA, Van Mieghem and his Network Architectures & Services (NAS) research group at Delft University of Technology have found a prospective solution that, theoretically, seems to have the best qualifications. A comprehensive simulation and test programme recently carried out together with American researchers demonstrated that the computation method also satisfied the high requirements under practical conditions.

Professor Piet Van Mieghem is convinced of one thing: Internet has to do much better! True, you can already listen to the radio through the World Wide Web, you can telephone and you can surf anywhere you want. And e-mails almost always reach their destination.
'However, no provider will guarantee that you can telephone flawlessly across the Internet or that you can hold an undisrupted video conference', according to Van Mieghem. 'There are simply no guarantees of quality.'
Segments of sounds or images may disappear, you may have to wait too long for an answer or the line just fails no matter how broadband your connection is. Yet the 'old-fashioned' alternative, the telephone, has been operating for decades with virtually none of those problems.
According to Van Mieghem, the design of the Internet needs a major overhaul. Providers will have to start offering quality guarantees for Internet connections and clients will have to start paying for them. Moreover, the control of the World Wide Web will become significantly more complex.

Secondary role The root of the problem is that the Internet was never built to transmit data within a particular time span.
'Basically, it was purely a data network in which time was initially a secondary concern', explains Van Mieghem.
Any file that is transmitted over the Internet, whether it is an e-mail, a webpage or an audio signal, is split up into packets, each of which is sent individually into the network. That network is an enormous, relatively unorganised system of interconnected computers. Each packet follows its own route before they are all neatly aligned again at the destination: the e-mail, the web page or the sound segment has arrived. How long it takes depends on the travelling time of the last packet, a fairly randomly determined parameter.
Moreover, data can also be lost, sometimes because of bad connections but mainly because the routers, busy junctions in the network, become temporarily overloaded and throw away incoming packets. These losses can be rectified by a system of confirmation of receipt and resending but that causes additional delays. Experts refer to this approach as "connectionless". Not because there is no connection between sender and recipient but because no fixed connection is reserved for a single data flow. Every packet finds its own way, almost haphazardly. The advantage is that the network is not particularly hierarchical and can easily cope with local malfunctions (the data simply take a detour).
In addition, the network can easily be extended without central control.
The counterpart of a connectionless network is the "connection-oriented" approach, exemplified for instance by the worldwide telephone network, which is composed of a strictly hierarchical pattern of local, national and international exchanges and connections. For a single telephone call to the United States a connection route is reserved for the duration of the conversation. Everything is designed to transport the data flow as flawlessly as possible, without losses and preferably as quickly as possible, because even a delay of more than a tenth of a second is a nuisance in a conversation.
According to Van Mieghem and many Internet developers like him, what is actually required is a sort of hybrid between the two extremes. This would involve a monitoring system that gives certain packets a higher priority, like registered mail and reserves bandwidth for it at a surcharge for the user.

"Quality of Service" (QoS) is the somewhat meaningless technical term that network specialists use for this approach of price and quality differentiation. As a matter of fact there was once a QoS-based data communication protocol, the "connection-oriented" ATM protocol for networks, which was developed by the International Telecommunication Union (ITU) standardization sector and ATM Forum, the foundation for broadcasting networks, in the late eighties and early nineties.
'It never took off, partly due to the Internet hype', says Van Mieghem with regret.

Problem child However, what is needed now is to improve the Internet machinery, says Van Mieghem. With the current state-of-the-art the international computer network will not be able to grow much further.
'The approach to Internet extensions has always been that if you just increase capacity everything will be all right', the professor says. 'That is basically how the Internet was built. But that won't work any longer.'
Routing, the switching of data flows over the Internet through busy exchanges or routers, is gradually starting to become a problem child. The Internet comprises some fifteen thousand subnetworks of major Internet providers, or 'Autonomous Systems' (AS). To find the best routes within their own subnetworks the providers use the OSPF (Open Shortest Path First) protocol, which is based on a computation method developed in 1959 by the Dutch mathematician Edsger Dijkstra (1930-2002). The major advantage of this 'Dijkstra algorithm' is its lightning speed, even in large networks. As the number of routers in the network increases the computation time to find the optimum path only increases a little more than proportionally.
'That is really awesome', according to Van Mieghem.
At a higher level, however, in powerful routers that make up the network backbone, the 'Border Gateway Protocol' is active. It does not determine the route using 'Dijkstra' but on the basis of arrangements made between Internet providers. Consequently, it does not necessarily result in the shortest route.
'Border Gateway Protocol is a really complex system with all kinds of updates, attenuators and other routing tricks to guide traffic properly', Van Mieghem says. 'Routing specialists are already discovering that BGP cannot be extended much further.'
The complexity and inefficiency are already causing major problems, for instance in the event of a BGP router failure. An alternative solution is a much more extensive type of monitoring, with which it is also possible at the highest level to obtain an overview of the entire network and to plan the best routes.
A QoS system with price and quality differentiation offers that possibility, according to Van Mieghem, because the system already needs an overview of the network for the division of the capacity among data packets with higher and lower privileges. In 1998, after much debate this led the Internet standards organisation IETF (Internet Engineering Task Force) to take over the 'Quality of Service' idea, which was then already ten years old, and this boosted research and interest in QoS systems. Experimental versions of such systems, for instance with twelve different priority classes, are being tested by dozens of research groups.

Infamous class Yet making the switch over is not so easy. From a computational point of view QoS routing is much more complicated than regular routing. The Dijkstra algorithm has just one criterion per subpath between two routers: the delay experienced by the data on the way. For efficient QoS routing, by contrast, several criteria have to be considered, including bandwidth, delay and the percentage of packets lost.
By contrast with 'Dijkstra', searching for a route that fully satisfies several criteria is an 'NP-complete problem', according to Dr Fernando Kuipers, who was until recently a doctoral student in Van Mieghem's group. In other words, QoS routing belongs to an infamous class of computation problems for which in some cases the computation time increases exponentially with the scale of the problem, in this case the size of the network. Quite simply, for some network structures every possible route must be computed, and with every node the number of possible paths may increase drastically by the number of nodes in the network. A network with several dozens of nodes may easily have billions times billions of possible routes and it will take even the most powerful computers days or weeks to compute them all. Such waiting periods are no option for the Internet.

Schematic overview of a small part of the global Internet. Each of these AS (Autonomous System) ovals represents a network of an Internet service provider (ISP) as seen from the Amsterdam Internet Exchange. The Internet comprises some 15,000 autonomous systems all over the world.

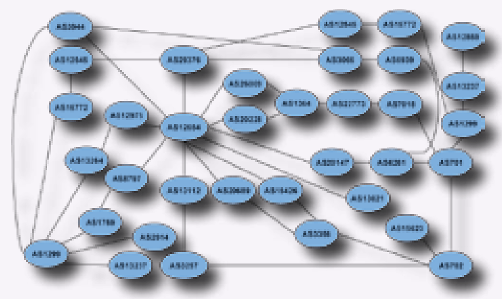

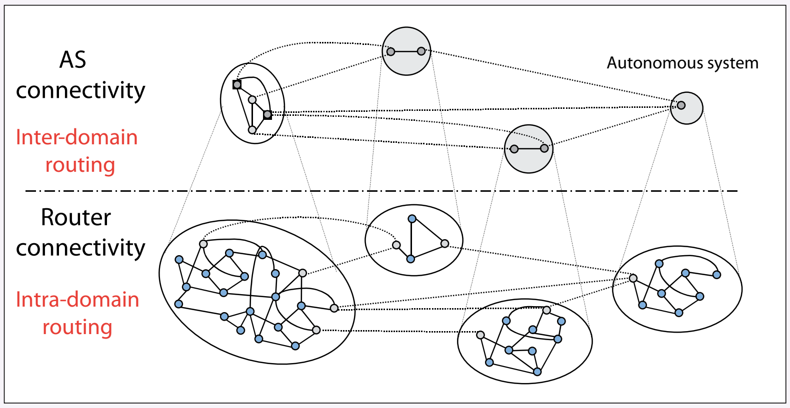

The Internet actually consists of two hierarchical layers. At the top is the AS layer used for traffic between Internet providers (known as inter-domain routing). Below it is the IP layer, where communication between users of the same Internet provider takes place (known as intra-domain routing).

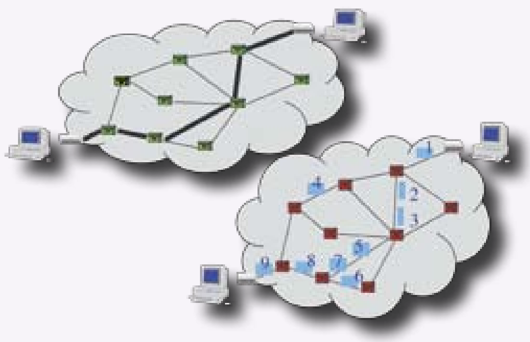

There are two ways to get a packet from A to B. Telephone is an example of the 'connection-oriented approach', which is jargon for the situation where all packets of a conversation follow the same path from A to B. 'Connectionless' transmission basically means that every packet of a data flow may follow a different path from A to B. The Internet is a classic example of 'connectionless' transmission

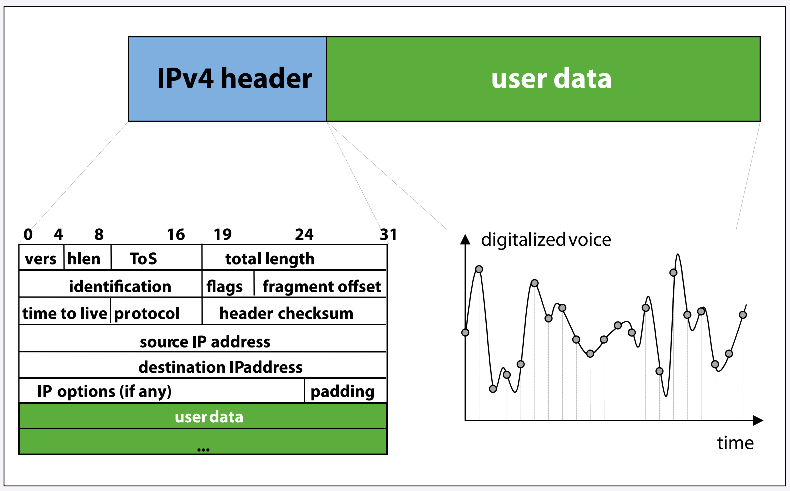

All information transmitted via the Internet is digitised. The data is sent as a data packet. A control section is placed in front containing among other things the IP address of sender and recipient, the total length of the packet and the protocol used to generate the data (the information).

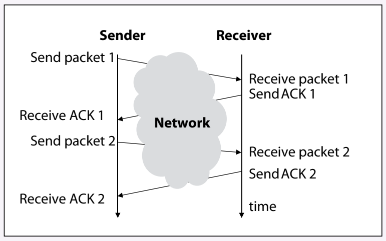

When data packets are transmitted over the 'connectionless' Internet some packets may not reach their destination. To ensure that IP packets that have not arrived are still sent TCP (Transmission Control Protocol) uses a technique in which the receipt of the packet is confirmed with an 'ACK' message.

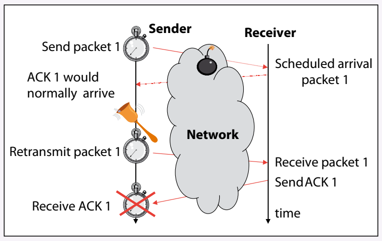

When source and recipient decide to communicate through the TCP protocol the source expects a confirmation (ACK message) for every data packet sent. However, the source will not wait forever for such an ACK. If the packet is lost in the network the source will wait for a certain period. If it has not received any message within the anticipated period the lost packet is resent. This process is repeated until a confirmation has been received.

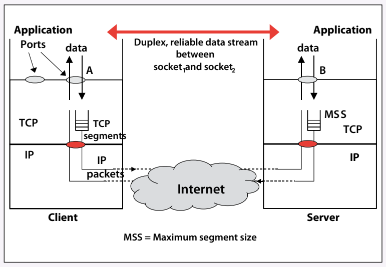

Communication between two e-mail programs on different computers over the Internet follows a layered structure. The top layer is known as the application layer where the programs generate information. On the layer below it TCP divides the information into segments. The size of the segments depends on the available capacity of the connection A-B. The TCP segments are sent as IP packets over the Internet, whereupon TCP regroups them into a single correct information flow at the recipient's end.

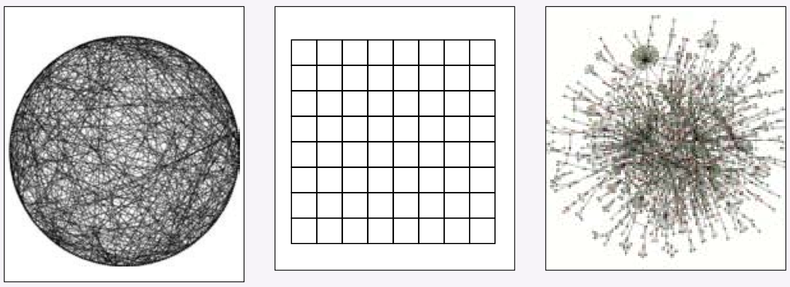

There are various types of network topologies. Two extreme types are the random topology and the lattice topology. The lattice topologies are highly regular, while their counterpart, the random topologies, are irregular. The topology of the Internet probably lies somewhere between these two extremes. The third topology, which besides for instance the protein structure probably also includes the Internet, is referred to by physicists as 'scale-free'.

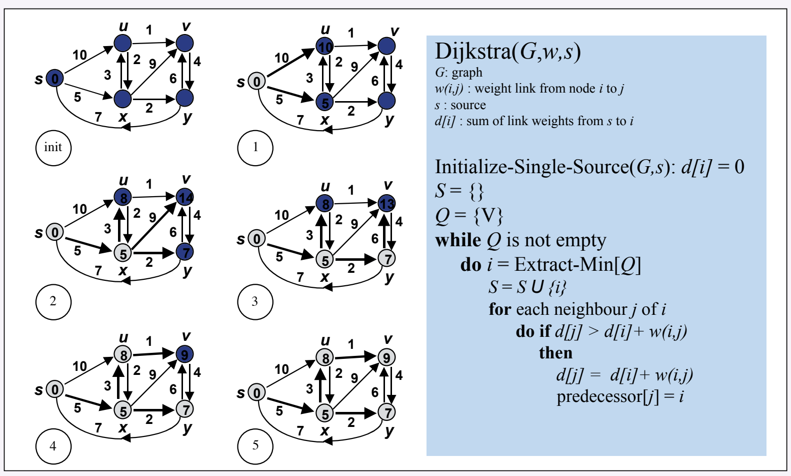

The Dijkstra algorithm finds the shortest paths from the starting point (s) to all other nodes in the network. The distance from s to all other nodes d[i] is initially infinite, except for s itself: d[s]= 0. Each time Dijkstra chooses the node i for which d[i] is smallest and checks whether it can improve the path to the neighbours of i by going through i. Once a node has been chosen it is never chosen a second time. So when all nodes have been chosen 'Dijkstra' stops and the shortest paths have been found.

'This NP completeness tends to deter researchers', Kuipers says. 'Their reasoning is that if you cannot solve it exactly easily, you just have to make a stab at it. That leads to 'heuristics', computation methods based on reasonable sounding rules of thumb. For example, preferably look for short routes through the network even though a detour may cause less delay. This means that heuristic methods do not guarantee that the optimum path is found, but at least the computation time remains limited (even though that cannot always be guaranteed either).'

Van Mieghem: 'There is currently a boom in publications in which people compare their heuristic with a different heuristic which happens to be a little better or a little worse.'

Van Mieghem himself, who was until 1998 employed by telecom company Alcatel in Antwerp, initially also followed that strategy and developed the heuristic algorithm TAMCRA, which stands for 'Tunable Accuracy Multiple Constraints Routing Algorithm'. Unfortunately, Alcatel decided not to develop the idea further at the time.'

SAMCRA Van Mieghem then started in Delft where he and his doctoral student Kuipers decided to have another shot at finding a non-heuristic, exact solution. This resulted in a new variant: SAMCRA (Self Adapting Multiple Constraints Routing Algorithm). SAMCRA is an algorithm that does guarantee it will find the path that best meets with all criteria (or reports with certainty that there is no such path). Van Mieghem acknowledges that in theory this can take a very long time.

'But NP-complete means explosive computation time in the worst case.'

In practice, there were never any problems of extremely long computation times with realistic network structures, the Delft researchers noticed. To substantiate that observation, Kuipers carried out a comprehensive test programme in which he tested SAMCRA for a large variety of network structure types. Eventually he found four conditions that all have to be met for NP-complete behaviour to emerge. Most of these conditions were not very realistic and are unlikely to occur in a regular computer network. One such condition was a clearly negative correlation between, for instance, delay and the rate of packet loss. However, connections between two routers in fact tend to lose more packets when delays increase, Kuipers explains. The researchers considered it highly unrealistic that the four conditions would coincide.

'The beauty of it is that you can conclude that realistic networks can still be computed properly with an exact algorithm', Van Mieghem states. 'We have developed more algorithms, but SAMCRA is our thoroughbred.'

Together with researchers at the University of Arizona and the University of Texas, Austin, they compared SAMCRA with a large number of heuristic algorithms by simulating a series of realistic networks on a computer.

'In every case SAMCRA outperformed the others', Kuipers reports.

The two cannot however prove that SAMCRA will never behave as NP-complete in realistic networks, Van Mieghem admits. 'If you could do that, you would not be far from winning a major prize in mathematics.'

Kuipers recently obtained a doctorate *cum laude* for his work on the algorithm, which the network researchers are trying to disseminate among Internet and network experts.

'The code is on our site', says Van Mieghem invitingly. 'Another nice aspect', he suggests, 'is that SAMCRA can also be applied to other routing problems. Route planning for vehicles, for instance: while Dijkstra's algorithm can simply find the shortest route, SAMCRA could also take into account delays caused by congestion or traffic lights or, if you want, the number of restaurants along the road.'

Kuipers was once approached by a robotics researcher who wanted to use SAMCRA for an independent robot. The robot worked on the basis of information from various cameras. When integrating the various video images problems similar to those of routing cropped up. Other applications, for instance for logistic problems or when searching large databases, are also conceivable.

Van Mieghem: 'I welcome any suggestions'.

Business model 'Broad application of SAMCRA and the quality differentiation of QoS would immediately help Internet providers to improve their 'business model', the professor says. 'Today most Internet providers charge a flat fee for which you can surf as much as you want (up to a certain limit).

They sell the bandwidth under an obligation to provide best efforts, not with a quality guarantee.' But to keep the revenues coming in once everyone is on the Internet and there is no longer any market growth you are going to need price differentiation, Van Mieghem predicts.

'In America there have been studies into what types of infrastructures are profitable in the long run', Van Mieghem explains. 'They showed that the postal mail service and railroads, for instance, barely pay their way. Only the telephone network, in this case America's AT&T, generated a return after sixty years.' According to Van Mieghem, apart from the monopoly position enjoyed by the telephone companies this was also due to the 'connection-oriented' organisation of the telephone network where you pay more for more extensive facilities (international phone calls are more expensive than local ones) but get guaranteed quality for your money.

SAMCRA could be a first step in that direction for the Internet too, Van Mieghem thinks, on the way to the ultimate goal of the 'Holy Grail'. 'When we have reached that point you will be able to get the quality of connection you need for any application at any time and anywhere and at a reasonable cost.' 'Who needs all those quality gradations and guarantees? Banks that want guarantees that their money transactions are safe', the professor suggests, 'but also companies that want to hold video conferences or to telephone over the Internet. I know plenty of people who say: Voice over IP (internet telephony – editor) works fine for me. But what if 20 million people want to use it rather than 200,000 as at present, and not only in the evening? Then of course there are other applications that you and I cannot even think of yet. E-mailing complete movie libraries, real-time graphic applications in 3D, who can say? History has shown that if you offer more bandwidth people will think of something to fill it.'

For more information on this subject, please contact Piet Van Mieghem, phone (015) 278 2397, e-mail p.vanmieghem@ewi.tudelft.nl or
Fernando Kuipers, phone (015)278 1347, e-mail f.a.kuipers@ewi.tudelft.nl

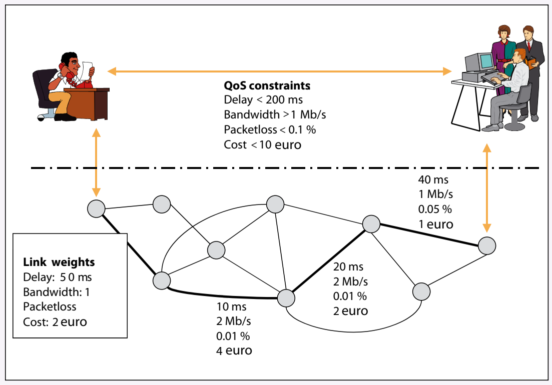

The QoS routing problem involves finding paths that comply with several requirements (e.g. delays, bandwidth, probability of packet loss and costs). For instance, VoIP (internet phoning) can only perform properly if the delay is lower than 200 ms, the bandwidth is not less than 110 Kb/s, the packet loss probability is lower than 1% and costs are as low as possible. The network may be composed of several connections (glass, copper, satellite) each with its own characteristics. The quality properties of such a path are determined by the weights of the various links within that path. The trick with QoS routing is to find the path that remains within the predefined QoS requirements.

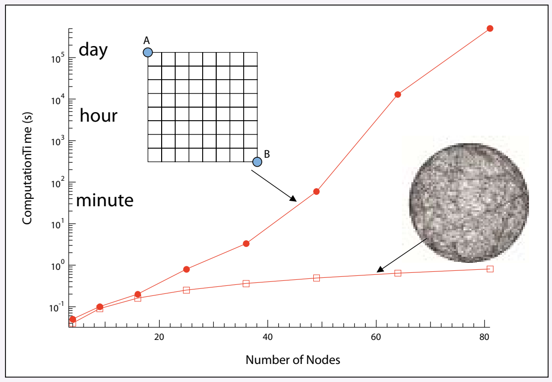

NP-complete problems are problems for which in the worst case the number of computation steps and the computation time can very easily get out of hand. The QoS routing problem is such a problem, but fortunately the computation time appears to be very short in practice.

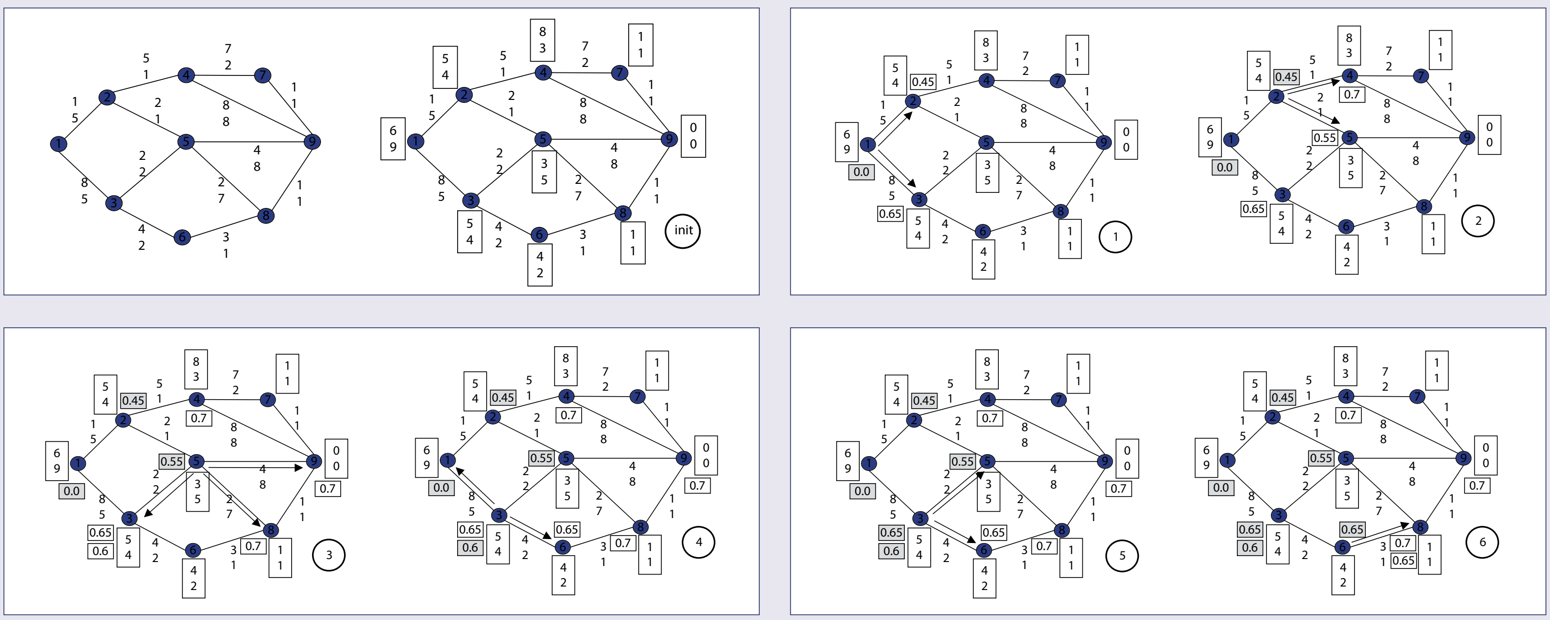

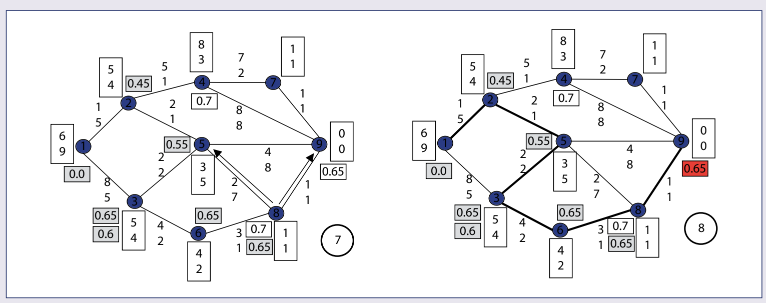

An illustration of SAMCRA. The network is characterised by two parameters: delay and costs. The example shows a search for the path between node 1 and node 9 that complies with the requirements, in this case set at 20,20. That means the delay must be less than 20 units and the costs must not exceed € 20. In the initialisation stage the Dijkstra algorithm is used to find the lower limit for the distance from each node to the destination. These lower limits bounds are shown in the rectangles. SAMCRA (just like Dijkstra) then starts to choose the path with the shortest length and extends this path to the neighbours under certain conditions. SAMCRA's length function is non-linear and it is shown in the small rectangles. Once a path has been chosen it is coloured grey. Step 3 shows that various paths can be established for a node. SAMCRA stops (step 8) when the path with the shortest distance to the destination has been found.

Digital technology extends the boundaries of panoramic photography

Modern Mesdag fits the

Panorama image of the path along the cemetry, with the TU Delft Assembly Hall across the water, February 2004. The distorted horizon makes it more difficult to see that this is in fact a 360° picture (the left edge connects to the right edge).

BY JOOST VAN KASTEREN

In 1881 Hendrik Willem Mesdag painted his famous Panorama of the fishing village of Scheveningen. The panorama covers a cylinder 14 metres high with a circumference of 120 metres. It is in fact an early form of virtual reality, giving the viewer the illusion of participating in life on the seashore. Like a modern Mesdag, Aldo Hoeben, part-time lecturer at the ID-StudioLab of the faculty of Industrial Design, creates digital panoramas that draw the viewer into events such as the funeral of Queen Juliana or a visit to the jungle of a zoo. The techniques used and developed by Hoeben for his panoramas can help designers to visualise their ideas and concepts.

At the ID-StudioLab at the TU Delft faculty of Industrial Design, one of the research subjects is 'inspiration engineering'; the development of tools and techniques to support designers in the development of ideas and concepts during the early stages of the design process, including the communication with other designers as well as clients. For this purpose, Hoeben is designing a digital sketchbook that, like a paper pad, enables designers to develop and evaluate their ideas through internal and external dialogues (see frame text). On the face of it, developing a digital sketchbook would seem a far cry from making panoramic views of a zoo or the Market Square in Delft, but there is more to it than meets the eye, according to Hoeben. "In both cases I am trying to get the imitation inside the computer to match what people are used to seeing in real life. For instance, the most-used sketches in the digital sketchbook get dog-eared like pages of a paper pad. In my panoramic view I am trying to use software to increase the level of apparent reality, for example by facilitating natural movements and through gradual transitions from light to dark.
In both cases the main thing is not to spoil the user's expectations through technical flaws," says Hoeben.
"Panoramic photographs are an interesting aid for industrial designers and students," he adds. "They can be used to give an impression of the surroundings in which future products are to be used. To develop such an image most people still use a collage or 'mood board'. Another option is to make a video collage as used in the TRI system, which was also developed at the ID-StudioLab (see Delft Outlook 2004-3). A panoramic photograph is a bit of both. It is relatively easy to make, though it takes more time than a video collage does, and it offers a lot of spatiality and conveys the atmosphere."

Crooked perspective Besides his part-time job as a TU Delft lecturer, Hoeben works in his own company, fieldOfView, where he develops and markets techniques for panorama photography. Photographers used to think that panoramic photographs were a hobby that got out of hand. Special

Digital sketchbook

He has forsaken his pencil, pen, and paper. Aldo Hoeben now uses a digital sketchbook, based on a Tablet PC with a touch screen. This makes him a walking prototype of the designer of the future, entrusting his ideas and concepts to a digitiser rather than putting them on paper.
Hoeben: "The computer already plays a major role in the design process, both for industrial designers, architects and graphical designers, but mostly at the end of the design process, when ideas are converted into working drawings and 3-D images. During the initial stages, when designers are still searching, they prefer to put their ideas onto paper."
Tablet PCs offer new possibilities, in particular now that their price has started to go down. "But," Hoeben says, "They will succeed only if we manage to tune into the designer's methods during the initial stages."
This means for example that to start with the digital sketchbook display is empty, without any tool bars or other aids that purport to make things easier but in fact do nothing but force another person's aesthetics onto the user. There is no array of virtual pens either, just a quartet of scribbles in the lower right corner made by the user himself — virtually speaking — using a black and a red fine liner, a blue marker, and a grey pencil.
Hoeben: "We teach our students to start a drawing by making a scribbles at the bottom to see whether a pen or marker still works. This is mimicked by the digital sketchbook. All you have to do is touch a squiggle to change the stylus from a fine liner into a marker or vice versa."

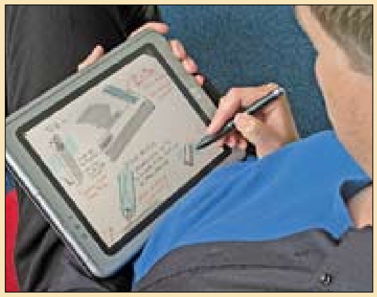

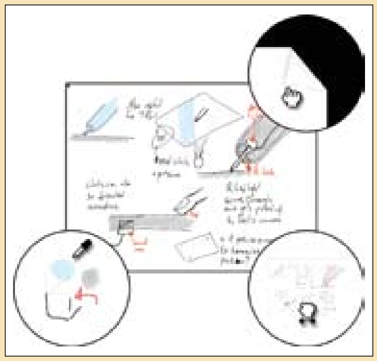

To preserve the largest possible area of the display for sketching, some of the action has been relegated to the corners of the screen. In the bottom left-hand corner of each page is a set of scribbles that the designer can use as a palette to select various types of pen.
The sketchbook pages can be leafed through by clicking in the lower right-hand corner and dragging the styles across a thumbnail image of the current page. To make a page easier to find, the top right-hand corner of the page can even be dog-eared by the designer.

Dog-eared

A key property of the paper sketchbook is that you can get an idea of the design history just by turning its pages, going from the first tentative jottings to more detailed sketches later on. Of course, turning pages on an electronic display is always a bit cumbersome. You may have to switch from page to page, or you may be using a scroll bar that does not show which page you end up on until you release it. In the lower right-hand corner of Hoeben's digital sketchbook there is a thumbnail view of a sketch page on which you can just about make out the individual sketches.
Hoeben: "If the sketches and drawings are your own, you can recognise them from the thumbnail."
Clicking and dragging the thumbnail makes the previously filled pages of the sketchbook appear in a fixed order. Pages that are checked more often remain slightly longer in view than the other pages. After a while, these pages even become dog-eared.
Hoeben: "These are all simple aids to make the digital sketchbook conform to the expectations of the designer."

world into his computer

cameras were used to create panoramic views of large groups of people but the technique was costly and cumbersome. Of course, you can always resort to sticking together separate photographs to form a panorama, but the results are strictly for the family album only. The problem is that the lines of perspective of such composite panoramas are all wrong. Unless the subject is a wide landscape view, the perspective effect will always be slightly off.
Thanks to the advent of the digital camera and image-processing software the panorama photograph has made a full comeback. Not only do the new techniques make it much easier to create panorama photographs with the correct perspective, but additional software has been developed that enables you to walk through the panorama, as it were. You could visit an estate agent's web site to walk through a house without disturbing the occupants. Just as we can experience late nineteenth century Scheveningen by visiting Mesdag's Panorama in The Hague, so modern digital panorama photography now enables us to feel the atmosphere of a wintery day in Holland, or a pavement in Washington, DC.

Tracing To make his panoramic photographs Aldo Hoeben uses a fish-eye lens on his digital camera. This type of lens captures a hemispherical view, covering just over 180 degrees in any direction.
A photographer pointing the camera straight ahead will also see the toes of his shoes. To make a complete panorama, Hoeben does an about-turn with the camera and takes four pictures within the space of ten seconds to record a full view of the surroundings. Using special software the four overlapping photographs are then combined into a single image. A special feature of the software enables the user to correct the distortion introduced by the fish-eye perspective, so crooked towers can be straightened out. Adding several fish-eye photographs together produces what is known as an equi-rectangular panorama in which the surroundings are recorded not only 360 degrees around in the horizontal plane, but also 180 degrees from top to bottom. The result is like projecting a view of the entire surroundings onto the inside

Fish-eye view of the interior of the Mesdag Panorama in The Hague. To increase the panorama's illusion, the edges of the painting are hidden from view.
An important element is the so-called 'faux terrain', a physical stretch of landscape that practically seamlessly blends in with the painted panorama. Visitors view the Mesdag Panorama from a beach pavilion surrounded by a section of imitation dunescape.

Panorama painters soon discovered that the standard tricks for creating a (central) perspective did not work on the curved surface of a panorama. To ensure that the paintings appeared natural, they developed a new trick. The artist sat with his head inside a large glass cylinder onto the surface of which he traced the surrounding view. The glass cylinder used by Mesdag is on display at the Mesdag Panorama Museum.

A fish-eye lens is an extreme wide-angle lens, with a horizontal field of view of just over 180°. Unlike a normal wide-angle lens, it introduces a characteristic distortion that shows straight lines as curves.

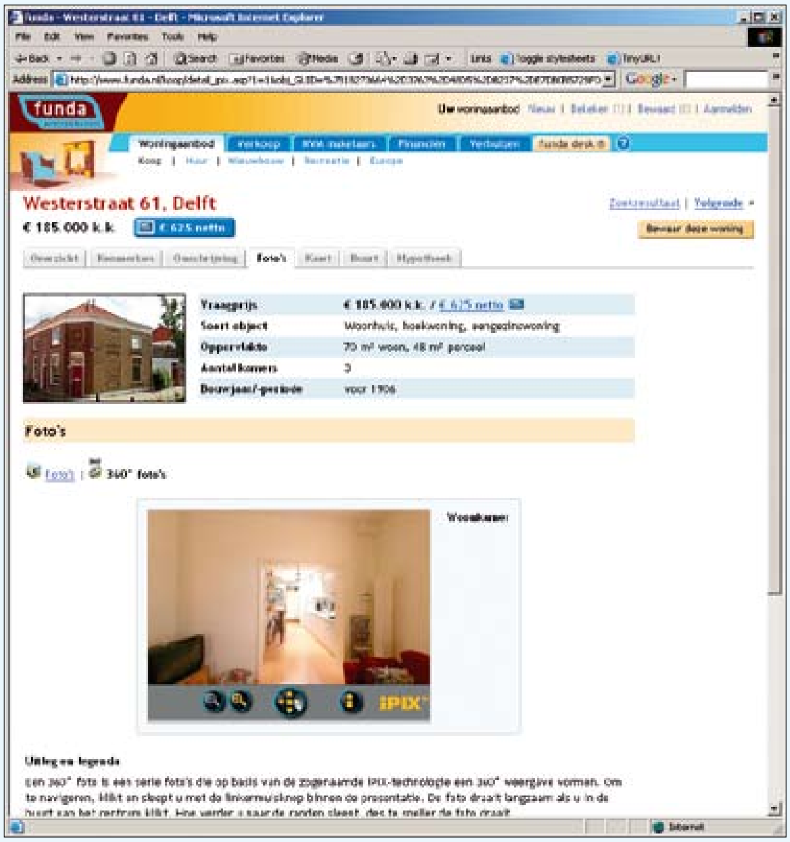

Some estate agents already use interactive panoramas on their web sites to show the interiors of the houses they have on offer. Whereas separate photographs on an estate agent's web site often give the idea that they show a single view of the house at its most advantageous, a panorama also reveals what is on the other side of the room.

When taking photographs for a panorama, it is vital to rotate the camera around the nodal point of the lens. Hoeben (shown here on The Binnenhof, the seat of the Dutch government in The Hague) uses a plumb bob instead of the usual tripod. It enables him to easily maintain a fixed position and height for each exposure.

The fish-eye images form sections of a sphere. By combining the separate fish-eye views — 4 in this case — a full view of the surroundings can be recorded. The overlap between the various photographs is used to fit the views seamlessly together. In spite of the fact that these pictures record the entire surroundings, the photographer himself is not included in the picture as he stood behind the camera for each exposure.

of a sphere that is then reworked to create a flat surface. In fact, a similar process was used to create the Mesdag Panorama. Just like other panorama painters of his day, Mesdag traced the contours of the view he wanted to paint onto the inside of a glass cylinder surrounding his head. The actual cylinder he used can still be seen in the Mesdag Panorama Museum in The Hague. Using this aid, Mesdag managed to get the perspective of a circular painting just right. This is not a trivial matter, since the lines of perspective, which normally extend as straight lines to the vanishing points on the horizon, are now curved along a cylinder or spherical projection. Once the contours had been sketched on the inside of the cylinder, they could be traced onto a long rectangular sheet of paper to produce a flat version.

Horizon to sinusoid

In addition to spheres and cylinders, other geometrical shapes can also be used as projection screens for images of the surroundings. Hoeben shows a photograph of the interior of a church in Dordrecht, in which the raw image has been projected onto a cube which was then unfolded to produce a cross. Another example is a photograph of the Scheveningen Fireworks Festival featuring a hyperbolic perspective.
In the resulting image of this seaside event, the full horizon has become a circle, with buildings, and the fireworks display sticking out of the almost planetoid scene. Cylindrical projection also offers plenty of scope. Standing on the north side of the Market Square in Delft, Hoeben took several photographs shortly before the arrival of the funeral procession of Queen Juliana. A cylindrical projection of the panorama produces an image in which both the town hall and the church can be seen from the front even though they are on opposite sides of the square. At first sight, it appears to be a normal picture. But the fact that the guard of honour is arranged in a horseshoe shape rather than in a straight line makes the viewer aware of the considerable distortion involved. A very interesting effect can be created by putting the virtual cylinder at an angle rather than vertically. This turns the horizon into a sinusoid. To illustrate the effect Hoeben shows a picture of a dog he took last winter near the cemetery behind the University's Assembly Hall.
Hoeben: "The little dog got a bit lost in the original picture, so I tilted the axis of the cylinder at an angle that put the animal in the upper half of the photograph, the centre of attention. As an additional effect, the horizon became distorted into a sinusoid."
The photograph also illustrates the time dimension involved. A woman approaching along the path from the right can be seen walking into the distance to the left in the same picture.
Hoeben: "Panorama photography enables you to capture not just three, but even four dimensions on a flat surface."

Patents cause stagnation

The software used to convert the different projections was created by German mathematician Helmut Dersch. Hoeben was involved in its development on the user side. Sadly, developments in the panorama software sector are threatened by stagnation as a result of a patent granted to an American company, IPIX. Like many U.S. patents, its scope is rather wide, and even though the patent in itself does not apply to Europe (where software patents did not become possible until very recently), the company has a policy of aggressively tackling any party they suspect of infringing it.
Hoeben: "As a result of this aggressive policy, few companies are interested in improving the software for creating panoramic views. Consequently, a number of panorama photographers all over the world have decided to develop the required programs themselves, either as open source software in large groups, or commercially in small, elusive companies.
Hoeben himself is one such developer. His fieldOfView company offers the SPi-V Engine, short for Shockwave Panorama Viewer, and known as Spiffy. It enables the user to use a mouse to click on the image and then drag the image from left to right or up and down (see also www.fieldofview.nl). It feels like standing at the viewpoint and looking all around you, just like in the Mesdag panorama. The SPi-V Engine is based on Macromedia's Shockwave technology, and according to Hoeben it does not infringe the IPIX patent being entirely based on techniques that had been published before the patent application was submitted.
These techniques were, and still are, used to display three-dimensional

objects in games like Doom, or feature films like Terminator II. The drawback of classic 3-D objects is they either lack realism or require huge amounts of processing power.
Hoeben: “The SPi-V Engine on the other hand enables you to create a very realistic environment in real time based on photographs in which you can look all around you. What I am trying to do is add elements that will enhance the realistic viewing experience.”

Soft braking One of these elements is an inertia effect in movements. In other panorama viewers the image halts rather abruptly as soon as you stop moving the mouse. Hoeben developed software that continues the movement for a fraction of a second. According to Hoeben, this feels much more natural, because in the physical world a person’s neck does not stop instantly either. To test his theory, he asked a number of students to look around in a panorama photograph with and without the inertia effect.
“When people view the image without the inertia effect, nobody notices anything wrong. The same applies when they subsequently view the image with the inertia applied. However, if we reverse things to test the inertia effect first, and then view the image without it, they do notice a difference. Some even felt the lack of the inertia effect in their necks. This shows that a built-in inertia effect is a closer match for the way people actually move, and consequently, it makes viewing the image a more natural experience.”
Another addition is the transition from light to dark sections of the image. The full dynamic range, which is the difference between the brightest white and the darkest black, is easily covered by the human eye, which gradually adapts to changing lighting conditions in a matter of 1-1.5 seconds. A panorama photograph may contain a very wide range of values from light to dark. Hoeben shows a photograph of a cave in the Jordanian desert where the sole source of light is the entrance of the cave. The aperture is brightly lit, but the inside of the cave remains too dark to distinguish objects. Hoeben has developed software that gradually adapts the image in the same way that the human eye gradually adapts to changing lighting conditions. The interior now lights up to reveal a number of objects, including a person in a corner.

Seattle Besides enhancing existing objects in the panorama photograph, virtual objects can be added to enhance the image. A special example of this technique has the Seattle Space Needle as its subject. Standing on the tower’s observation deck Hoeben shot photographs of the view using his fish-eye lens. He then added an image of the Space Needle itself, or rather, a computer-generated image he created using classic three-dimensional graphic modelling techniques. Using a technique developed by Paul Debevec at the Institute for Creative Technology of the University of Southern California, Hoeben adjusted the lighting of the virtual tower to match the ambient lighting. The technique uses each pixel in the panorama as a light source with a certain colour and intensity. Not only do you see sunlight shining on the tower just like it does on the surrounding area, but you also get reflected light illuminating it, from a blue sky for example. The result is extremely convincing.
“So much so,” Hoeben says, “that many people think I actually flew around the Space Needle in a helicopter. They even ask me how I managed to get permission, since aircraft have been banned from airspace close to tall buildings ever since 9/11.”

Animation Apart from these almost imperceptible additions, the inclusion of a small animation may help to liven up the image, according to Hoeben. He shows a panoramic photograph of a so-called boardwalk exhaust, where a central heating plant releases its exhaust gases over a pavement somewhere in Washington, DC. Using a particle generator, a piece of software that generates moving particles, he can show smoke coming from the pipe, and even vary the direction and force of the wind. A small moving element brings the entire image to life. Another example is a movie of a talking woman that has been added to a panorama photograph of a building interior.
Hoeben: “Panorama photographs are often used to show what a room looks like, for example a hotel room or the living room of a house an estate agent is trying to sell. However attractive the room may be, the image still looks empty if it does not include people. A little movie like this, which in this case runs as

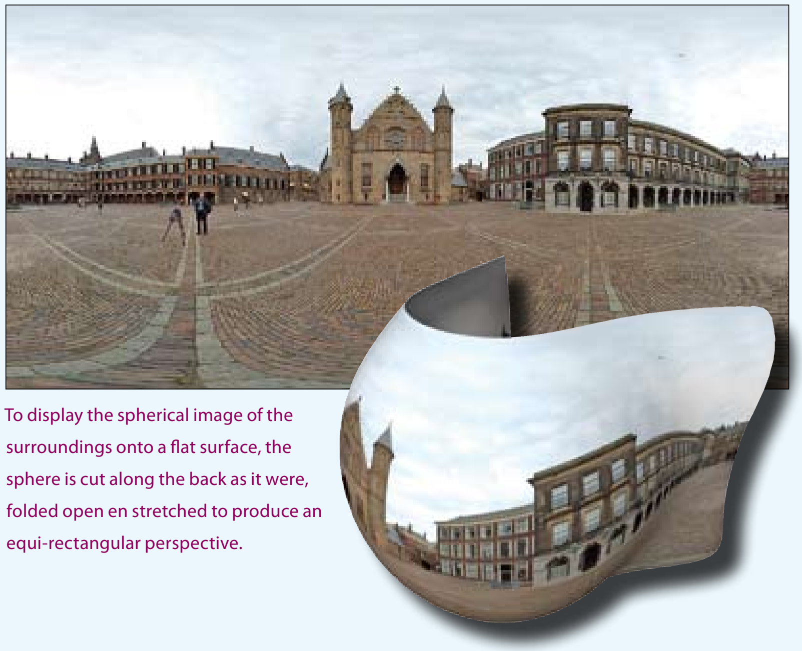

To display the spherical image of the surroundings onto a flat surface, the sphere is cut along the back as it were, folded open en stretched to produce an equi-rectangular perspective.

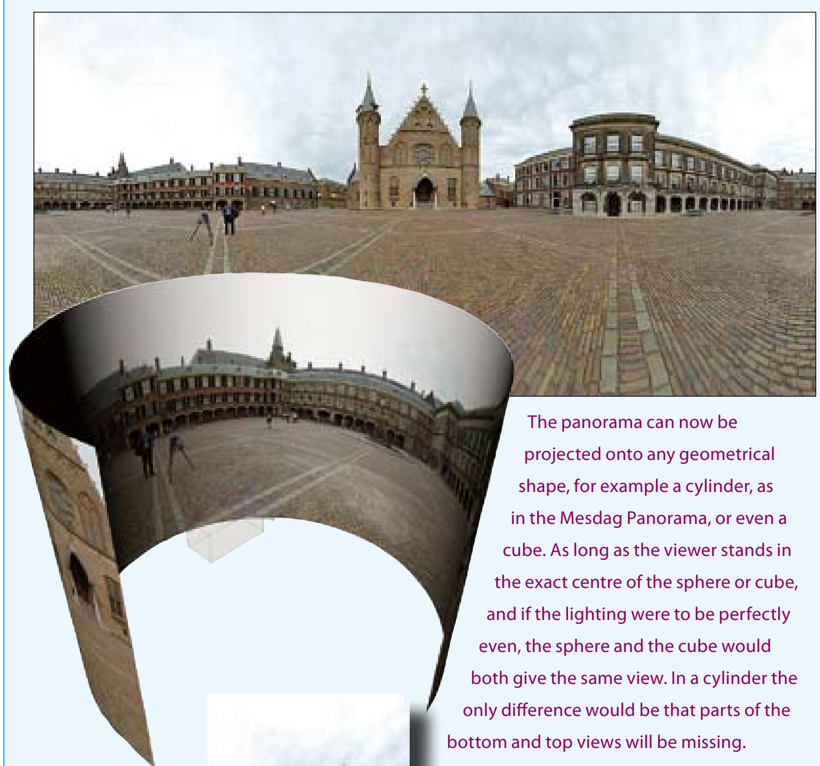

The panorama can now be projected onto any geometrical shape, for example a cylinder, as in the Mesdag Panorama, or even a cube. As long as the viewer stands in the exact centre of the sphere or cube, and if the lighting were to be perfectly even, the sphere and the cube would both give the same view. In a cylinder the only difference would be that parts of the bottom and top views will be missing.

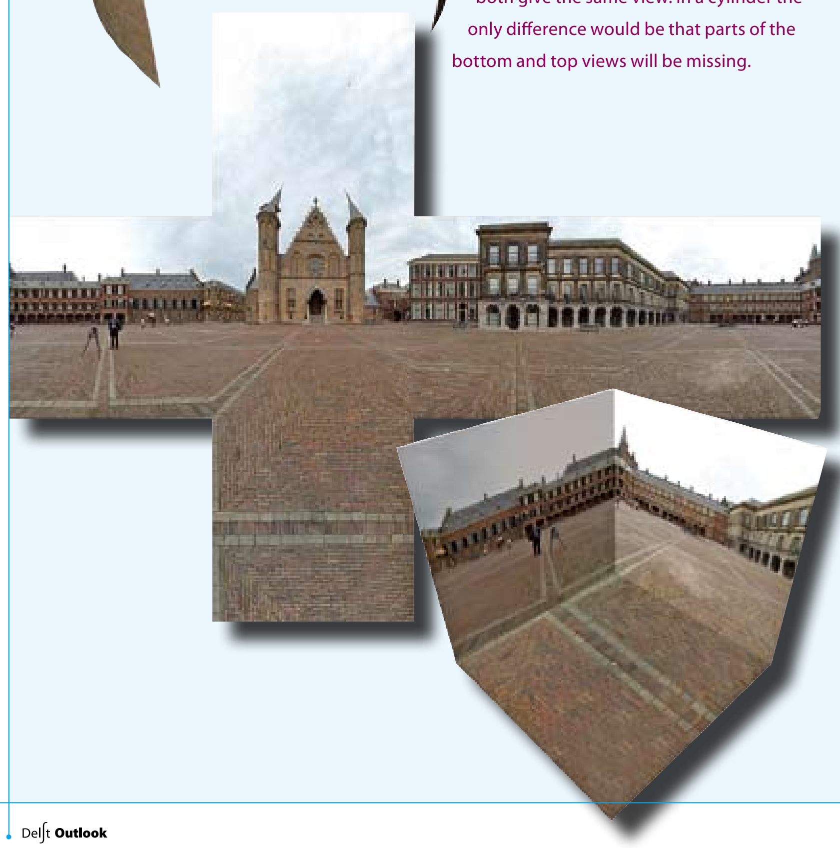

Any normal wide-angle lens can also be used to create a 360° × 180° panorama, although it requires more exposures. The first row of 16 photographs together form a cylinder. By photographing several rows, rotating the camera around its focal point, the entire surroundings can be recorded. In this example, a total of 52 exposures is required.

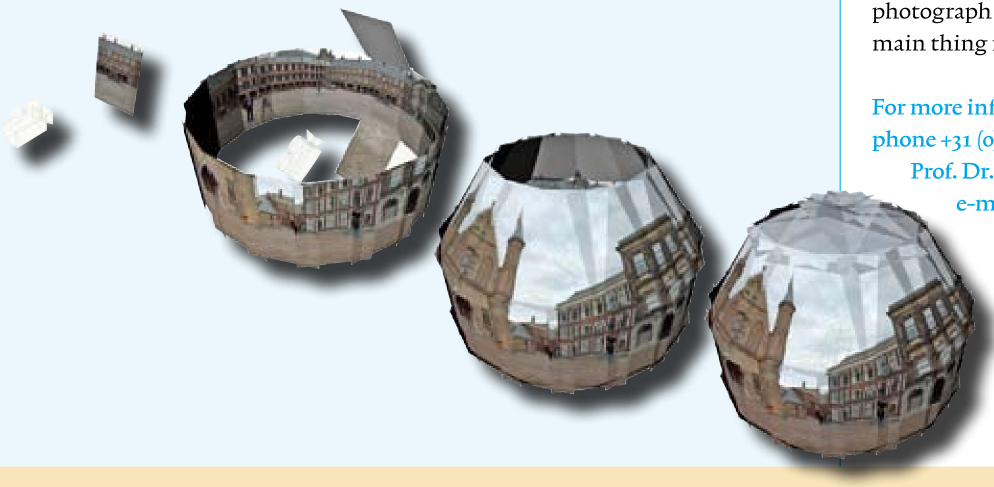

a permanent loop, adds authenticity to the image."
Authenticity is what it is all about.
Hoeben: "People should not be distracted by the technology, however beautiful or innovative. Whether you are looking around inside a panorama photograph or, as a designer, jotting down ideas in a digital sketchbook, the main thing is that you can really experience the space and objects."

For more information, please contact Ir. Aldo Hoeben, phone +31 (0)15 278 1714, e-mail a.hoeben@io.tudelft.nl or Prof. Dr. Ir. Pieter Jan Stappers, phone +31 (0)15 278 5202, e-mail p.j.stappers@io.tudelft.nl.

Relevant web sites:
www.fieldofview.nl
www.fieldofview.nl/events/outlook

The panorama Hoeben shot at the funeral of Queen Juliana in Delft shows how a panorama need not be a wide view of a landscape. In this case, the view of the crowd accurately expresses the elated, though at the same time subdued, mood of the occasion.

For web sites as used by estate agents, Hoeben came up with additional features for panoramas. In this case he has added some video footage of a woman to an interior view of a large house. She could be an estate agent explaining the layout or the history of the building to the virtual visitor. Note that the woman in the picture never actually entered the building; she was edited in later using digital processing software.

In his panorama of the Scheveningen fireworks festival Hoeben converted the spherical perspective into a hyperbolic perspective in which the virtual camera points vertically downwards. Lenses that can produce such a hyperbolic perspective do not actually exist.

Fotograph: Greg Downing, www.gregdowning.com

Most panorama photographers regard it as a challenge to achieve a good exposure throughout a scene. In the upper image the exterior view has the correct lighting conditions, whereas the interior of the cave is far too dark. On the lower of the two pictures, the lighting inside the cave is perfect, but the outside view is overexposed. Hoeben's software dynamically creates mixed lighting conditions depending on the direction of view in the panorama. The resulting lighting matches that observed by the human eye, which constantly adapts to differences in lighting levels.

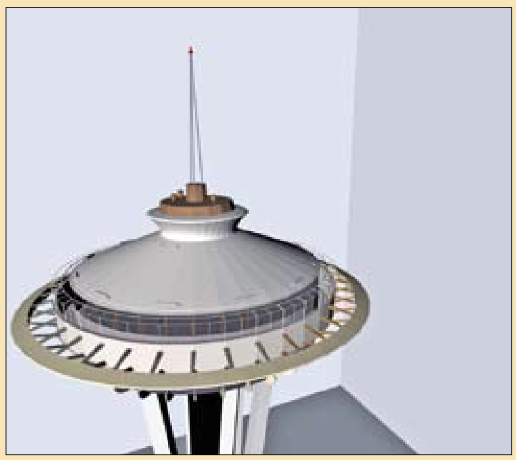

For the Space Needle demo Hoeben took photographs from the observation deck of the 180 metre high tower in Seattle. These enabled him to create a panorama of the area surrounding the Space Needle, although it did not include the tower itself. He then reconstructed a view of the tower using traditional 3D software and photographs he found on the Internet. The panorama not only provided a photo realistic background, it also helped to create the correct lighting conditions for the tower reconstruction. If — as in the upper view — the panorama and the lighting effects are removed, the tower becomes instantly recognisable as a product of traditional computer graphics software. The magic is in the combination of the panorama and the effect of the ambient lighting on the reconstructed tower.

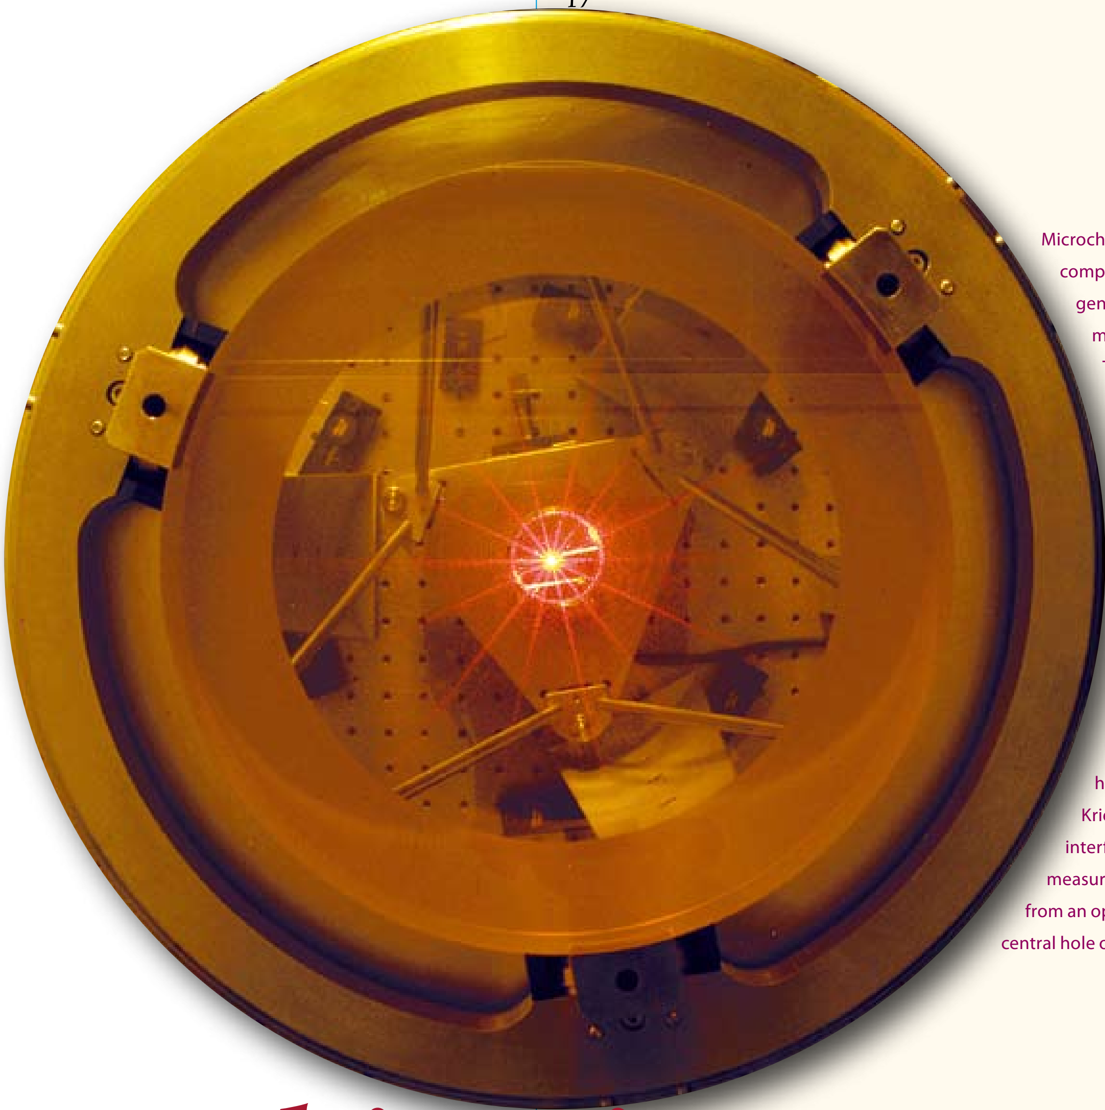

Microchips are getting more complicated by the day and future generations will feature even more intricate circuit patterns. The lithographic process that is used to project these circuits onto silicon currently still requires lenses. After 2010 however, the feature density on the chips will have grown so far that extreme ultraviolet light with a wavelength of 13nm will have to be used and the lenses will have to be replaced by mirrors. To measure these with the high accuracies required, Luke Krieg developed an absolute interferometer. This top view of his measurement set-up shows the light from an optical fibre shining through the central hole of a mirror being tested.

Lithographic mirrors
measured with sub-nanometre accuracy

The absolute interferometer for extreme ultraviolet optics

Future generations of microchips will probably be produced using extreme ultraviolet light with a wavelength of thirteen nanometres. Optical lenses will no longer be suitable for the manufacturing process because they absorb the light at such short wavelengths. Mirrors will have to be used instead. Currently the main problem is the required precision to accurately determine the curvature of the mirror down to the nanometre. Researchers at TU Delft have built a new measurement system that promises to achieve just such a precision: an absolute interferometer. This new measurement system offers many advantages over competing methods..

by Bennie Mols

The microchip industry is trying to create ever-smaller structures on its products. Smaller and smaller circuits mean faster processors, more memory capacity and more efficient computers. Current production methods use light to write structures on microchips. As the structures get smaller, the wavelength of the light must be reduced also. According to the rules of optics the smallest structural dimension is related to the wavelength of the light used to create it. Microchip manufacturers already use light with a wavelength of 193 nanometres (visible light lies in the region of between 400 and 700 nanometre).

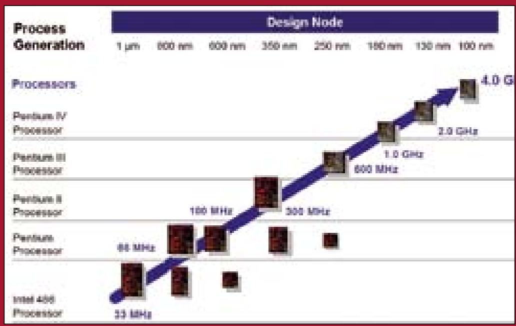

ILLUSTRATIE: CARL ZEISS SMT

According to Moore's Law the capacity of microchips will double every eighteen months. Although the American Gordon Moore, initially opted for a period of one year in 1965, and later changed his mind to two years, his final choice of eighteen months has yet to be proven wrong. The first lithographic processes used 436 nm light to create electrical connections 5 microns (1/200 mm) wide. The current generation of equipment uses 193 nm light and can make connections only 100 nanometres wide. Within a matter of years lithography is expected to use 157 nanometre light, although sceptics think this step will be skipped to pass directly to 13 nanometre technology.

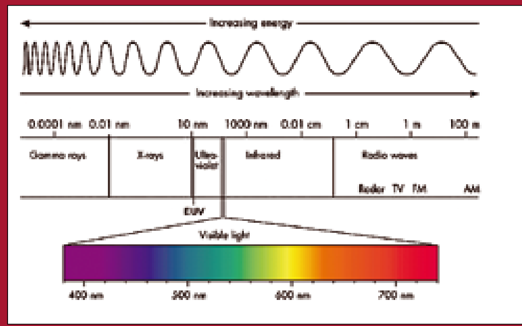

Electromagnetic radiation affects our daily lives in many ways, in both positive and negative senses. Radio waves enable us to watch television, use mobile phones and cook in a microwave oven. Infrared offers us heat, night-vision goggles and remote control devices. Visible light, which is responsible for the wide range of colours we can see, covers only a small part of the electromagnetic spectrum. Although generally speaking the high-energy radiation outside the visible range, i.e. UV, extreme UV, X-rays and gamma rays, can be harmful to living organisms, it is also used in medical applications, astronomy and the semiconductor industry. The current lithography systems are limited to light of wavelengths longer than those of extreme UV, because this light would be absorbed by the lenses in the equipment. The only way to write microchips with this type of light is to use mirrors.

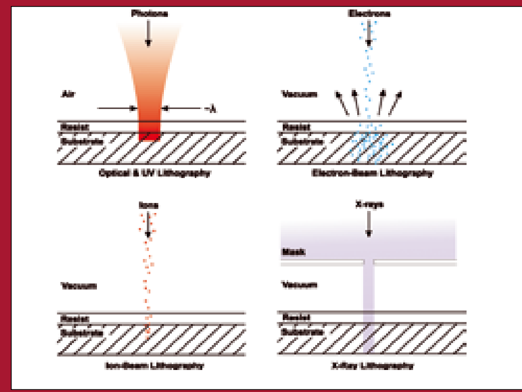

Four different methods of writing patterns onto silicon chips. Although visible-light lithography and UV lithography are limited by the wavelength of the light being used, these techniques can be readily scaled and they are easy and fast to use. Even finer patterns can be made using electron or ion beam lithography, but for the time being these techniques remain slow, expensive and prone to teething problems. X-ray lithography also has its drawbacks, in particular regarding construction and maintenance. The mask needs to be practically in contact with the surface being written on and no practical photo resist has yet been discovered for use with X-rays.

(PHOTOGRAPH: COURTESY OF SANDIA NATIONAL LABORATORIES / WWW.SANDIA.GOV)

Test set-up of an extreme UV exposure system developed by Sandia National Laboratories. The heavy stainless steel structure is needed to ensure sufficient stability of the writing system and structural integrity for working in vacuum.

This drawing shows how the extreme UV (purple) travels from mirror to mirror through the exposure system.

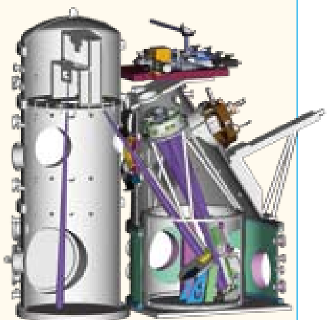

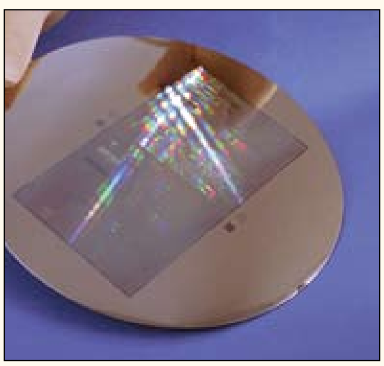

(PHOTOGRAPH: COURTESY OF SANDIA NATIONAL LABORATORIES / WWW.SANDIA.GOV)

Photographs of an extreme UV mask. It contains the patterns that will eventually be projected onto thousands of chip wafers.

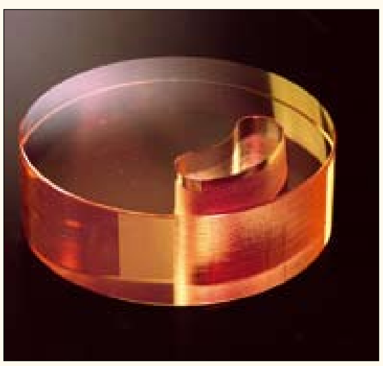

(PHOTOGRAPH: COURTESY OF SANDIA NATIONAL LABORATORIES / WWW.SANDIA.GOV)

Example of an extreme UV mirror substrate, made from Zerodur. This type of glass, with its characteristic yellow colour, is designed to minimise the changes in shape and size due to temperature changes. The reflective coating is deposited onto the practically transparent material in several layers.

Microchip manufacturers are planning to create microchip structures smaller than 35 nanometres by the beginning of the next decade. This can only be done using extreme ultraviolet light. The problem with this is that the manufacturers will no longer be able to use optical lenses, as at such short wavelengths these absorb all the light, so no light would pass through them. The solution is to use mirrors, because such short-wavelength light can still be reflected. The wavelength of the preferred light is 13 nanometre, an extremely short wavelength, strongly absorbs at these wavelengths, but since the reflective coatings made of silicon are very thin, this absorption has only a small effect. The main complication when using mirrors is the need for them to be very accurately shaped, down to one-tenth of a nanometre, which is about the diameter of a single hydrogen atom. Is there a measuring principle that can offer such precision, is not too expensive and can be used in a conventional optical workshop? Quite a challenge, but even so researchers at the TU Delft have managed to come up with just such a measuring method. Continuing from the preliminary theoretical work of doctorate student René Klaver, another doctorate student, Luke Krieg managed to build a measurement system that promises to achieve the required accuracy. Note the 'promises' in the previous sentence. Krieg has not quite been successful yet, but he is convinced that he will be so very soon. He works at the Optics section of Professor Joseph Braat at the faculty of Applied Physics at TU Delft. The research is supported by optical systems manufacturer Carl Zeiss, lithographic equipment manufacturer ASML, the research establishment TNO-TPD and the Netherlands Technology Foundation STW.

Absolute interferometry The system devised and constructed by Krieg is an absolute interferometer. The interferometer uses two interfering beams of light to measure the almost spherical shape of the concave mirrors. One beam falls on the mirror surface, where it is reflected before hitting a sensor, while the other beam of light reaches the sensor unobstructed. The word "absolute" refers to the fact that the only difference between the two beams of light is that one hits the mirror while the other one does not - there are no additional optical elements present. All competing interferometers are non-absolute, in that they use one or several intermediate optical elements, such as lenses, fibre surfaces, gratings and prisms. Because these elements are never entirely perfect, they introduce additional errors of their own which destroy the accuracy of the measurements.
"Our idea was to create an interferometer that did not use any other optical

elements to measure the mirror surface," explains Krieg about the measuring set-up. "It was to be a completely new type of interferometer"
On paper, the final setup appears very simple, but in practice, things are much more complicated. The interferometer can be divided up into four major sub-components: The light source, the interferometer frame, the sensor and finally an algorithm to turn the complicated readings from the sensor into the shape of the mirror.
Krieg: "Each component was a separate challenge. I spent about a year on each of them."
First however, he had to deal with the difficult concept of absolute interferometry. "The ultra-precise measurement of optical surfaces is essentially a chicken-and-egg problem", Krieg explains. "The shape of the surfaces is measured using a verified length standard - the wavelength of light from a stabilized Helium Neon laser. There is no question whatsoever about the accuracy of this length standard. The problem is in the path that this light takes through the measurement setup. If we place any optical elements, like lenses, in our light beam, we disturb its path. The only way to make an accurate measurement is to know this path exactly, for which we need to know the shape of the intermediate optics with the same accuracy as we need to know the shape of our mirror. To measure these we would need exactly the type of setup we're trying to build in the first place!"

Imperfect optics

And then there is another problem. The shape of the optical elements can change with time.
Krieg: "If we were to measure the shape today and then calibrate the entire system for that shape, we would not be able to repeat the measurement next week. The temperature might be slightly different, which would affect the shape of the lenses, so we would have to recalibrate all over again. Also, a surface may change after some time due to 'wear'. If all optics were perfect, it would be simpler to use optical elements in the measuring set-up, but they are not.
The only way to get an absolute measurement is to make sure the beam of light does not have to pass through any optical element. This makes for tough design specifications."
Krieg shows the mirror he used for his preliminary measurements, now safely packed away in a box. It has a diameter of twenty centimetres and it is transparent and yellowish in colour. The price of such a prototype is at least a million euros. Although the shape of the concave mirror appears to be perfectly spherical, it is ever so slightly off, in the order of a few micrometers, which is intentional. The mirror consists of a glasslike substrate that expands and contracts as little as possible with changing temperatures. The substrate will eventually be coated with approximately one hundred extremely thin layers to

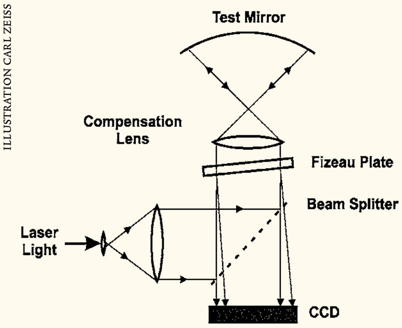

ILLUSTRATION CARL ZEISS

The Fizeau interferometer currently being used to measure extreme UV substrates contains a number of optical elements, including lenses, prisms and splitters. Paradoxically, all of these elements must first be tested with an accuracy equal to that of the instrument itself, creating a chicken-and-egg problem. Even once the shape of each optical element is known, they would still have to be measured over and over again to take into account any changes over time.

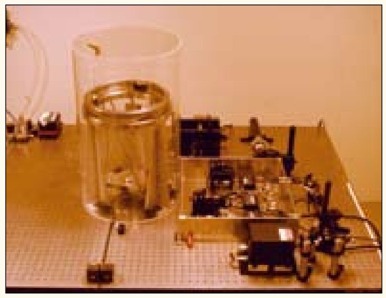

In broad terms, the prototype of the absolute interferometer developed by the Optics Research Group at Delft University of Technology consists of a hexapod frame (to the left) and a light source, a box with a complex optical system (to the right). Both are connected by means of two optical fibres. During actual measurements the cylinder is filled with helium to minimise any residual turbulence effects caused by minute temperature variations.

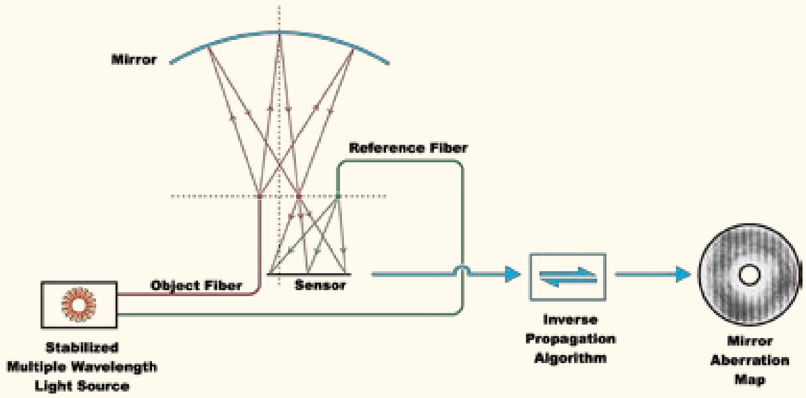

The light exits the optical block through a pair of glass fibres. The light from one of the fibres is reflected by the mirror and focused between the two optical fibres before hitting the sensor, where it interferes with the light coming from the other fibre. The continuously changing interference pattern is recorded by the sensor. The measured data is processed using an Inverse Propagation Algorithm (IPA) developed by Krieg, which enables the shape of the mirror to be calculated down to sub-nanometre level.

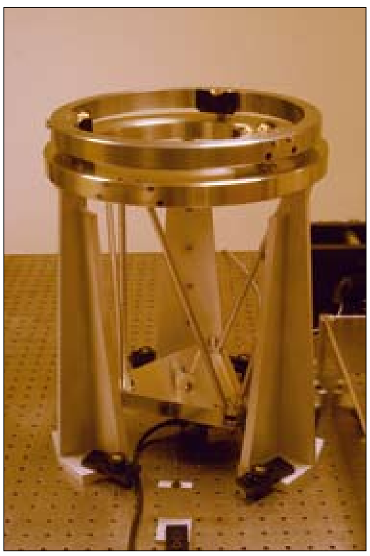

The task of the interferometer's suspension is to ensure that the relative position of the optical fibres, the sensor and the mirror remain constant within a few millionths of a millimetre. The structure is made from a special alloy with a thermal expansion coefficient one-tenth that of stainless steel. The mirror being measured (not shown) is placed on the ring at the top of the instrument, which is designed to ensure that the mirror always settles in exactly the same position every time it is inserted.

The light of each of the two lasers has to be focused to a spot of only a few microns before it can be effectively coupled into the optical fibre.

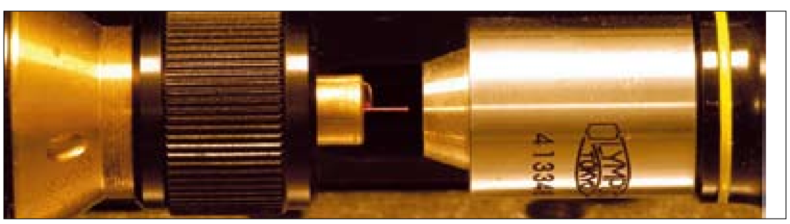

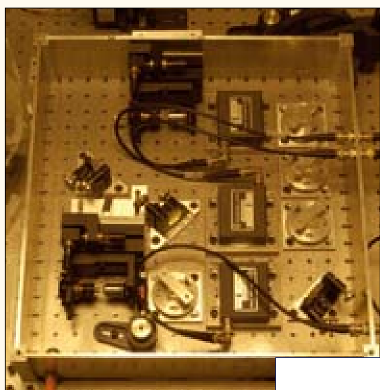

Krieg's flexible light source can work with several different optical techniques, so that it can be used to measure with different sensors. The beams from three lasers are sent through an array of mirrors, splitters, modulators, wave plates and optical fibres before they reach the mirror under test. A feedback system adjusts the wavelength of the tuneable laser to lock it relative to the stabilised laser.

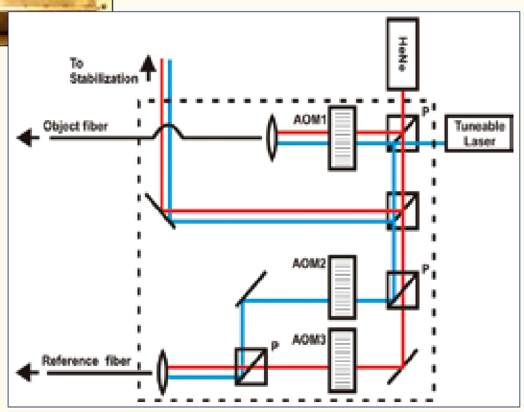

This schematic diagram of the light source shows the path travelled by the laser beams. The source produces lights of both 633 and 637 nanometres.

Precision mirrors for astronomers, physicists and biologists


ILLUSTRATION: NASA

The Chandra space telescope launched in 1999 by NASA to detect novas, pulsars and other high-energy sources already uses mirrors for the X-ray range. Improved measuring methods will help to greatly enhance the resolution of such telescopes.

Microchip manufacturers are not the only people looking for extremely accurate curved mirrors. Astronomers measuring in the extreme ultraviolet light range (usually between 5 and 30 nanometres) also need them. Extreme ultraviolet radiation is produced by gases with temperatures of at least one million degrees. The Sun's corona and the direct vicinity of black holes can produce this type of radiation. Extremely short-wavelength light is also released when stars explode. Physicists analysing the conditions inside a plasma (a mix of free electrons and ions of the type produced inside nuclear fusion reactors) also have to work with extremely short-wavelength light for which optical lenses are useless. Biologists studying minute structures inside living cells also use increasingly short wavelengths. Under normal conditions, water absorbs very strongly at these short wavelengths, which would result in very low contrasts. However, there is a very small wavelength region where no absorption occurs in water, known as the water window. Biologists are now looking for ways of harnessing this light. Normal optical elements can no longer be used and mirror-based microscopes will have to be developed.

As in lithography, shorter wavelengths mean higher resolutions for microscopy. In 'in-vitro' biological applications very short wavelengths can be a problem however, since water absorbs most of this light. There is only a small wavelength range around 3 nanometres, the water window, in which the absorption by water is less than that by materials of interest such as proteins. Again, mirrors with extreme accuracy will be needed to achieve the desired resolution.

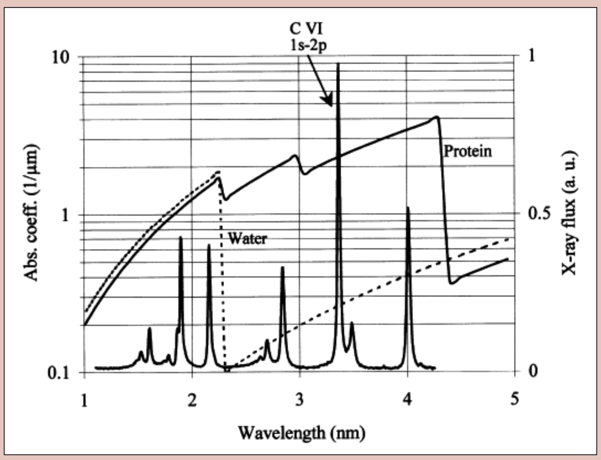

DRAWING: INTERNET

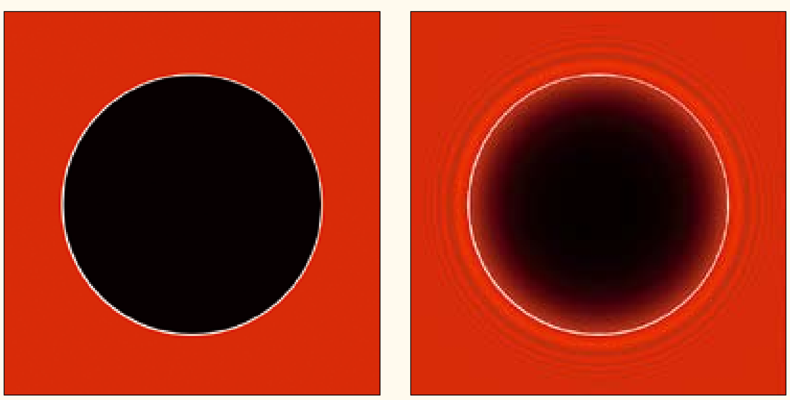
Many mirrors feature a hole or discontinuities to satisfy the requirements of optical designers. At the edges of these areas, diffraction occurs, affecting the accuracy of the measurements.

make the mirror as reflective as possible to extreme ultraviolet light. The final result is a balance between good reflective properties and the lowest possible absorption.

Trick & helium The trick used by Krieg to achieve absolute interference relies on the special properties of optical fibres. An optical fibre is placed near the centre of curvature of the mirror. The fibre has to be close to the centre of curvature, but not too close, or the light will be reflected back onto the tip of the fibre, which must be avoided. The fibre core conducting the beam of light measures only three micrometres in diameter. The laser light entering the fibre at the light source exits at the fibre tip as a perfect spherical wave, propagating radially in all directions like an expanding sphere. The spherical wave then strikes the mirror where it is reflected onto a light-sensitive sensor. On the other side of the centre of curvature of the mirror is the tip of a reference fibre. This also emits a spherical wave, but this one hits the sensor directly. The reflected light and the reference light interfere with each other, and the sensor registers the interference pattern.

"In fact, we are comparing light reflected by our mirror with light that passes unobstructed," Krieg says. "Any difference is entirely the result of the shape of the mirror. This means we only have to calibrate the system once, unlike other methods."

Since air also acts as an optical element with a certain refractive index, an absolute interferometer cannot operate in air. One solution would be to put the set-up in a vacuum, but this is far from simple. It is much easier to cover the measurement system with a cylinder and flush the air out, to fill the cylinder with helium gas instead.

Krieg, lifting the cylinder off the optical bench: "The refractive index of helium is practically the same as that of a vacuum. It is much easier to supply a bottle of helium gas with our system than to have users supply their own vacuum system."

Lines, rings & chaos On a computer display, Krieg shows the resulting interference pattern. It looks much more complicated than the usual interference patterns of lines and rings. In fact, it's chaos. Krieg also knows that the number of interference lines exceeds the number of pixels in the sensor measuring the interference pattern. This means that the sensor's resolution is too low to see the full interference pattern.

Krieg: "Any scientists working in interferometry when presented with a pattern like this will say it's no use to them. Nevertheless, we managed to solve the interpretation problem."

Krieg uses two interference patterns, produced by using light with two slightly different wavelengths. He starts with a standard helium-neon laser with a wavelength of 633 nanometres and then uses a 637 nanometre laser light source. Both are in the visible part of the spectrum and both produce an interference pattern, but the patterns differ slightly from each other.

"By subtracting one pattern from the other," Krieg says, "we get a new interference pattern, which we can interpret."

Active stabilisation No commercially available light source could produce an ultra stable wavelength different from but still close to 633 nanometres. So the research team had to make their own.

"We started with a standard tunable diode laser," Krieg explains, "the wavelength of which can be adjusted. However, we have actively stabilised it to produce only light of 637 nanometres."

The interference pattern produced between the beam reflected by the mirror and the direct reference beam hits a CCD sensor. The sensor currently used by Krieg is ten years old and lacks the required accuracy.

"As part of the same research project, the faculty of Electrical Engineering, Mathematics and Computer Science at Delft University is currently developing a completely new type of sensor that meets the specifications of our absolute interferometer. The sensor is being specially developed for our test system, but unfortunately it is not ready yet. The project started later because it proved difficult to find a suitable doctorate candidate."

In its current incarnation, the Delft interferometer can determine the shape of the mirror down to four nanometres, whereas the requirement is one-tenth of a nanometre.

"Even so," says Krieg, "the sensor we are currently using is the main bottleneck. The fact that we can get an accuracy of four nanometres using this old sensor is a good indication that the new sensor being developed by DIMES will enable us to achieve the one-tenth nanometre requirement. I am convinced we will get there, since all the other parts of the system satisfy the accuracy requirements."
Note that a dimensional accuracy of one-tenth of a nanometre does not mean that the position of each atom on the mirror's surface has been measured. It simply means that the overall shape of the mirror will be measured within one-tenth of a nanometre. The system does not look at small local deformations. Other methods already exist to measure those in a statistical sense.

Countless benefits In addition to the great benefit that the absolute interferometer does not introduce systemic errors caused by the effects of optical elements along the route, the system offers other benefits.
"Our system can handle much greater deviations from a perfect sphere than other interferometers can," says Krieg. "Other interferometers always have problems measuring deviations from a perfect sphere, and as the deviation increases, the problems increase . The lithographic industry alone will need many differently shaped mirrors. Some will deviate more from the perfect sphere, others less. Since we are using two different wavelengths, there are practically no restrictions. This increases the general applicability of the absolute interferometer. On top of that, our system could be extended to measure lenses as well."
There is yet another advantage. Many mirrors contain holes or consist of strange shapes, dictated by their application. Along the edges of a hole or any other discontinuity, light is diffracted. This effect introduces undesirable disruptions in the interference pattern, which seriously affect the accuracy of measurements at sub-nanometre levels.
Krieg: "Other methods measure only mirrors without sharp edges, and if the mirror does have them, they stay clear of the edges. Such methods cannot cope with diffraction. Our method is capable of compensating for diffraction."
First, the resulting interference pattern is reduced from a very chaotic pattern to something that lends itself to interpretation. Next, the disruptive diffraction effect is subtracted. Finally, the measurement system represents the shape of the mirror as a series of at least six hundred numbers, each of which provides information about which corrections are required at which location on the mirror. Each of these interpretation steps requires quite a bit of mathematics.

Competition The main competition for the absolute interferometer developed by the group at TU Delft is the Sommargren interferometer developed by the Lawrence Livermore National Laboratories in collaboration with the National Institute or Standards & Technology (NIST) in the U.S.
Krieg: "As early as 1999 Gary Sommargren said that he had an absolute interferometer, but that was not the case. His first models still used a lens right in front of the sensor, but that has now been removed. Even so, his system still

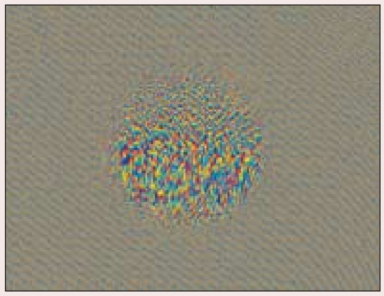

The measurement data obtained by Krieg's interferometer forms a chaotic interference pattern. The pattern contains more oscillations than the number of pixels in the measuring sensor. Hence additional information is needed to enable the pattern to be unambiguously interpreted nonetheless.. This information is supplied by a second interference pattern produced by a different wavelength.
The region in the centre of the image represents the area of the mirror with the hole in it. The pattern there is irregular because there is almost no light present. This area is not included in the measurement.

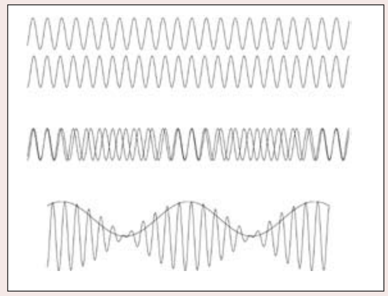

The interference patterns of two almost identical wavelengths differ only slightly, but when one is subtracted from the other, a pattern emerges that appears to be the result of light with a much longer, 'synthetic', wavelength.

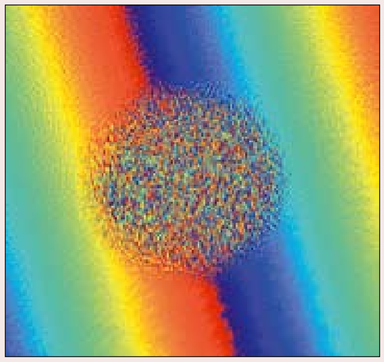

The synthetic interference pattern, measured using two different wavelengths, shows less variation than the pattern produced by a single wavelength.

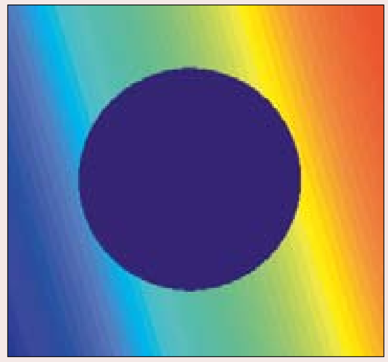

Physically speaking, this restored pattern is identical to the chaotic pattern produced by the first wavelength, but it uses the information contained in the more ordered synthetic interference pattern shown in the previous illustration. Only after this restoration procedure does the interference pattern provide sufficient information to calculate the shape of the mirror.

Competing nano lithographic methods for microchips of the future

In the race to establish the standard method for the production of the next generation of microchips, writing with extreme ultraviolet light (EUVL) is currently the most promising candidate. It is the subject of research projects in Europe as well as the U.S. and Japan. The International Semiconductors Roadmap describes EUVL as a potential solution for the future of microchip manufacture. Large companies such as ASML, Intel, IBM and Motorola are working on it, as are a number of large research establishments such as the Lawrence Livermore National Laboratory, the Berkeley National Laboratory and the Sandia National Laboratories. Even so it is far from certain that EUVL will become the next method of choice. Critics are of the opinion that the system requires a level of accuracy that is not feasible. One of the main obstacles to commercial use is the accuracy of the mirrors needed. The key question is whether the method will be sufficiently cost-effective. This has yet to be determined.
Another possibility, though often considered a second choice, is to write the patterns using electrons (Electron Beam Projection Lithography, EPL). The wavelength of electrons is one-ten thousandth of a nanometre, which is much less than the size of an atom. A beam of electrons can easily be focused into a point with a diameter of one nanometre. However, the benefits of the small wavelength are offset by several major problems. The production speed of the writing technique is far too low for use in mass production. Also, the resolution is limited by the dispersion of electrons in various steps of the writing process on the chip and includes dispersion by the chip's photosensitive layer.
Other existing methods include X-ray lithography and ion beam projection lithography. X-ray lithography works with a (very expensive) synchrotron that supplies X-rays with wavelengths of approximately one nanometre. A one-on-one mask then projects the radiation onto the chip wafer. To avoid high absorption rates, most of the installation has to be in a vacuum. X-ray masks are fragile and expensive to make and have remained a major obstacle to the commercial use of X-ray lithography. Ion beam projection lithography uses hydrogen or helium atoms. Complementary stencil masks are used to build up the image, which is then reduced by means of electrostatic lenses. The production of masks, the design of an ion imaging system, the fitting together of complementary masks and finding a suitable balance between sharpness and production speed remain the major problems with this method of writing microchips.

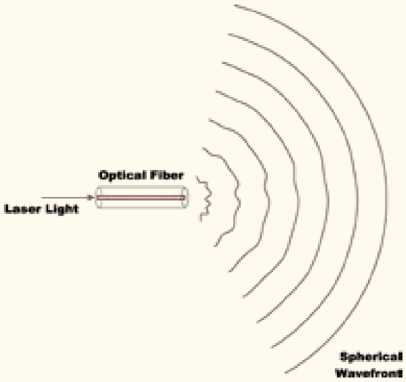

Ironically, the otherwise disruptive diffraction phenomenon also gives this interferometer its uncanny accuracy. Diffraction very rapidly changes the wave front of light coming from a small point (e.g. the tip of a glass fibre) into a perfectly spherical shape, even though the original wave front is irregular.

PHOTOGRAGH: LASER FOCUS WORLD

The Sommargren interferometer is also based on glass fibres, but unlike the interferometer developed by Krieg it uses optical elements between the fibres and the optics being tested.

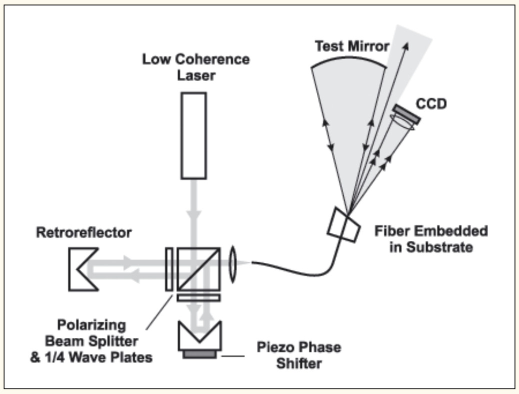

The light coming from the optical fibre in a Sommargren interferometer is reflected by the mirror being tested, back to the fibre, where the light is reflected under a different angle. It reaches the CCD, where it interferes with the rest of the light from the fibre. The reflecting surface of the fibre and the optics just before the CCD are optical elements. Just like the elements in a Fiezeau interferometer they must be accurately measured as they may become distorted over time. This is why the Sommargren interferometer cannot be called an absolute interferometer.

Lithography from Delft — the Mapper

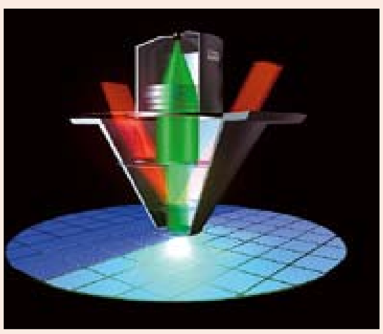

In July 2000 Prof. Dr. Ir. Pieter Kruit established the Mapper Lithography company, as a spin-off project from his academic work at TU Delft. Mapper is developing an alternative lithography method in addition to the methods currently receiving most of the world's attention: Extreme Ultraviolet Lithography (EUVL) and Electron Beam Projection Lithography (EPL). Mapper combines simultaneous writing using 13,000 electron beams with superfast optical data transport methods as used in the telecommunication industry. A single electron beam is split into 13,000 separate electron beams. Electrostatic lenses are then used to focus each of these beams onto the surface being written (the chip wafer). In a single scan, a wafer image field of 26×33 millimetres is exposed. The electron beams are switched on and off by 13,000 light beams, one for each electron beam. The light beams are generated by a data system containing information about what the required chip pattern looks like. Mapper's aim is to both write wafers rapidly (more than ten per hour) and achieve a high resolution (less than 45 nanometres). The method is currently in the developmental stage.

is not absolute. The fibre he uses to direct the light onto the mirror also serves as a reflective surface for the light returning from the mirror. Only then does it hit the sensor. The surface of the fibre can never be perfect and it is also prone to change with time. This makes his method less than absolute. He claims to be able to measure down to 0.25 nanometre, but that accuracy is not reproducible. If he were to repeat his measurement a week later, the outcome could be very different, even though the surface of the mirror under test hasn't changed. In addition, our system enables us to measure much larger mirrors in one go than the Sommargren interferometer can."

Other systems than Sommargren's use even more optics, which puts them that much farther from being absolute interferometers.

"Whether our absolute interferometer will eventually be developed for commercial use depends on many other factors," Krieg says. "The most important factor is the commercial feasibility of the proposed microchip production process."

Although the microchip industry remains the driving force behind the precision measurement systems for mirrors, the invention may also benefit biologists, physicists and astronomers. These scientists are increasingly involved in experiments requiring short-wavelength light and consequently mirrors with nano-precision shapes. The principle of the absolute interferometer developed at TU Delft is now protected by patent. Krieg will remain with the Optics section until the end of 2004. He hopes that the new sensor will become available before that time and that he will be able to use it to provide conclusive evidence that his absolute interferometer can measure the curvature of a mirror to an accuracy of one-tenth of a nanometre.

Krieg has no doubt that the new sensor design will bring the system up to scratch. With a bit of luck he will be able to demonstrate this himself, but if not, his successors will have their work cut out for them. Krieg has now completed his research and was awarded his doctorate on 30 November 2004.

For more information please contact
Dr. Luke Krieg,
phone +31 (0)15 278 2455,
e-mail m.l.krieg@tnw.tudelft.nl, or
Prof. Dr. Ir. Joseph Braat,
phone +31 (0)15 278 6146,
e-mail j.j.m.braat@tnw.tudelft.nl.

INS & OUTS

Recent publications by Delft University Press

Februari

The changing world of container Logistics
Binsbergen, A. van
ISBN 90-407-2474-1,
236 pages, € 47,50

Computational Support for the Selection of Energy Saving Building Components
Wilde, P. de
ISBN 90-407-2476-8, 240 pages, € 45,37

Maart

Waterlegislation in the Netherlands
Huisman, P.
ISBN 90-407-2463-6,
130 pages, € 27,50

Objects with perspective
Gerhardt, W.
ISBN 90-407-2464-4, 320 pages, € 60,00

Shear and Anchorage Behaviour of Fire Exposed Hollow Core Slabs
Fellinger, J.H.H.
ISBN 90-407-2482-2,
277 pages, € 45,37

Flying Glare
Beumler, T.
ISBN 90-407-2481-4,
488 pages, € 55,00

Mechanical behaviour and diffusion of gas during neutron irradiation of actinides in ceramic inert matrices
Neeft, E
ISBN 90-407-2475-X,
175 pages, € 33,92

Ápril

Megacorridors in North West Europe Investigating a new transnational planning concept
Zonneveld, W.
ISBN 90-407-2472-5,
104 pages, € 20,00

Multi-Objective Management of Saltwater Intrusion in Groundwater: Optimization under Uncertainty
Tran, T.
ISBN 90-407-2480-6,
208 pages, € 45,37

Inverse modelling of the rainfall-runoff relation A multi objective model calibration approach
Rientjes, T.
ISBN 90-407-2491-1,
411 pages, € 50,00

Knowledge-Based Design: Developing urban & regional design into a DUP Science
Klaasen, I.
ISBN 90-407-2479-2,
242 pages, € 37,00

Bubble size effect on the gas-lift technique
Guet, S.
ISBN 90-407-2492-X,
160 pages, € 33,92

Singe-molecule experiments on DNA with novel silicon nanostructures
Storm, A.
ISBN 90-407-2485-7,
134 pages, € 33,92

Slicing the earth: A layer stripping method employing a causality-based imaging condition
Poot, H.
ISBN 90-407-2490-3,
141 pages, € 33,92

Mei

Influence of Admixtures and Mixing Efficiency on the properties of self compacting concrete
Takada, K.
ISBN 90-407-2501-2,
249 pages, € 45,37

Prediction of Deterioration of concrete bridges
Gaal, G.
ISBN 90-407-2500-4,
162 pages, € 33,92

Integrating Reliability, Availability and Maintainability (RAM) in Conceptual Precess Design
Goel, H.
ISBN 90-407-2502-0,
157 pages, € 33,92

Stochastic inequality constrained closed-loop model predictive control
Hessem, D. van
ISBN 90-407-2489-x,
217 pages, € 45,37

Monitor Maritieme Arbeidsmarkt 2003
Aa, v.d. R.
ISBN 90-407-2495-4,
133 pages, € 25,00

A nanostructural study of Raney-type nickel catalysts
Devred, F.
ISBN 90-407-2503-9,
154 pages, € 33,92

Tensor product representations and special functions
Groenevelt,
ISBN 90-407-2504-7,
256 pages, € 45,37

Rules for Commercial Cruising Vessels
NML
ISBN 90-407-2493-8,
92 pages, 25,00

A novel tissue-equivalent proportional counter based on a gas electron multiplier
farahmand, M.
ISBN 90-407-2488-1,
160 pages, € 33,92

Effect of spatial variability on Maintenance and Repair Decisions for Concrete Structures
Li, Y.
ISBN 90-407-2505-5,
188 pages, € 33,92

Performance-based design of self-compacting fibre reinforced concrete
Grunewald, S.
ISBN 90-407-2487-3,
233 pages, € 45,37

Juni

Forming of Laminates
Jong, de T.W.
ISBN 90-407-2506-3,
288 pages, 45,37

Localization in dutch dune sand and organic clay
Cheng, X.
ISBN 90-407-2510-1,
108 pages, 33,92

High-rise housing in Europe
Turkington, R.
ISBN 90-407-2483-0,
300 pages, 49,95

Juli

Towards a methodology for large-sample prompt-gamma neutron-activation analysis
Degenaar, H.
ISBN 90-407-2509-8,
159 pages, € 33,92

Electronic instrumentation
Regtien, P.P.L.
ISBN 90-407-2525-X, 445 pages, € 32,50

Augustus

Sampling of Particulate Materials New theoretical approach
Geelhoed, B.
ISBN 90-407-2517-9,
208 pages, € 45,00

Parallel Hardware Architectures for the Life DUP Sciences
Epstein, A.
ISBN 90-407-2518-7,
160 pages, € 40,00

Modelling Passenger Flows in Public Transport Facilities
Daamen, W.
ISBN 90-407-2521-7,
404 pages, € 65,00

Optoelectronic Properties of Discotic Materials for device applications
Piris, J.
ISBN 90-407-2520-9,
159 pages, € 40,00

Quality of Service Routing in the Internet Theory, Complexity and Algorithms
Kuipers, F.
ISBN 90-407-2523-3,
243 pages, € 45,00

Photoinduced Charge Separation in Dye-Sencitized Films of Smooth and Nanocrystalline TiO2
Kroeze, J.E.
ISBN 90-407-2522-5,
198 pages, € 0,00

September

Atomistic simulation of Cu/Ta thin film deposition and other phenomena
Klaver, P.
ISBN 90-407-2526-8,
94 pages, € 25,00

Electrochemically fabricated nanometer-sized molecular junctions
Kervennic, Y.
ISBN 90-407-2513-6,
174 pages, € 40,00

Neutron Activation Analysis of inhomogeneous Large Samples
Baas, E.
ISBN 90-407-2528-4,
176 pages, € 40,00

Electron spin and charge in semiconductor quantum dots
Elzerman, J.
ISBN 90-407-2533-0,
160 pages, € 45,00

Oktober

Development of new storage phosphors for thermal neutron detection
Sidorenko, A.
ISBN 90-407-2540-3,
120 pages, € 35,00

Photogeneration and dynamics of charge carriers in the conjugated polymer poly(3-hexylthiophene)
Dicker, G.
ISBN 90-407-2542-X,
144 pages, € 35,00

Crustal rheology and postseismic deformation: modeling and application to the Apennines
Riva, R.
ISBN 90-407-2547-0,
125 pages, € 35,00

Instrumentation to handle thermal polarized neutron beams
Kraan, W.
ISBN 90-407-2527-6, 174 pages, € 40,00

New neutron-based isotopic analytical methods
Perego, R.
ISBN 90-407-2537-3,
189 pages, € 40,00

A World of Transport, Infrastructure and Logistics
Bovy, P.H.L.
ISBN 90-407-2555-1,
461 pages, € 60,00

Towards On-line Logistics: the LinC Interaction Modeling Language
Honig, J.
ISBN 90-407-2544-6,
327 pages, € 60,00

November

Antennas for multi-band RF Front-end modules
Boyle, K.
ISBN 90-407-2549-7,
339 pages, € 55,00

ICT Related Transformations in Latin American Metropolises
A.M. Fernandez-Maldonado
ISBN 90-407-2556-X,
316 pages, € 45,00

Policy integration in practice
Stead, D.
ISBN 90-407-2541-1
172 pages, € 24,95

Conformal Symmetries and constrained critical phenomena
Y. Deng
ISBN 90-407-2548-9,
209 pages, € 45,00

The Golgi: a transition point in membrane lipid composition and topology
Lisman, Q.
ISBN 90-407-2543-8,
112 pages, € 35,00

Spatial Flood Extent Modelling A performance-based comparison
Werner, M.
ISBN 90-407-2559-4,
188 pages, € 40,00

Filament Winding a unified approach
Koussios, S.
ISBN 90-407-2551-9,
400 pages, € 60,00

Present-day kinematics at the Azores-Gibraltar plate boundary as derived from GPS observations
Fernandes da Silva, R.
ISBN 90-407-2557-8,
212 pages, € 45,00

All publications of Delft University Press can be ordered directly through the TU Delft University Customer Service www.library.tudelft.nl, by phone +31 15 278 5678 or through international booksellers.

For purchases within the Netherlands, please add 6% value added tax.

WHO & WHERE

Delft University of Technology has eight faculties, each of which is engaged in education and research in one or more disciplines. The University was founded in 1842 by King William II. With 13.000 students, 2.800 scientific staff members and 2.500 technical and administrative employees, it is the largest university of technology in The Netherlands.

TU Delft
P.O. Box 5
NL-2600 AA Delft
the Netherlands
telephone +31-15 278 9111
telefax +31-15 278 6522

Disciplines

Aerospace Engineering
Kluyverweg 1
NL-2629 HS Delft
Telephone +31 15 278 1455

Applied Earth Sciences
Mijnbouwst raat 120
NL-2628 RX Delft
Telephone +31 15 278 1423

Applied Physics
Lorentzweg 1
NL-2628 CJ Delft
Telephone +31 15 278 5355

Architecture
Berlageweg 1
NL-2628 CR Delft
Telephone +31 15 278 4237

Chemical Technology & Bioprocess Technology
Julianalaan 136
NL-2628 BL Delft
Telephone +31 15 278 3659

Civil Engineering
Stevinweg 1
NL-2628 cn Delft
Telephone +31 15 278 5553

Electrical Engineering
Mekelweg 4
NL- 2628 CD Delft
Telephone +31 15 278 6260

Geodetic Engineering
Kluyverweg 1
NL-2629 HS Delft
Telephone +31 15 278 3289

Industrial Design Engineering
Landbergstraat 15
nl–2628 CE Delft
Telephone +31 15 278 3984

Life Science & Technology
Julianalaan 67
2628 BC Delft
Telephone +31 15 278 2413

Marine Technology
Mekelweg 2
NL-2628 CD Delft
Telephone +31 15 278 6595

Materials Science
Rotterdamseweg 137
NL-2628 AL Delft
Telephone +31 15 278 3659

Mechanical Engineering
Mekelweg 2
NL-2628 CD Delft
Telephone +31 15 278 6595

Computer Science
Mekelweg 4
NL- 2628 CD Delft
Telephone +31 15 278 4568

Applied Mathematics
Mekelweg 4
NL- 2628 CD Delft
Telephone +31 15 278 4568

Technology, Policy & Management
Jaffalaan 5
NL-2628 BX Delft
Telephone +31 15 278 7100

Multidisciplinary Centres

Adhesion Institute
Kluyverweg 3
NL-2629 HS Delft
Telephone +31 15 278 5353

Netherlands Institute for Metals Research (NIMR)
Rotterdamseweg 137
NL-2628 BL Delft
Telephone +31 15 278 2535
Fax +31 15 278 2591

Biotechnological Sciences Delft Leiden (BSDL)
Julianalaan 67
NL-2628 BC Delft
Telephone +31 15 278 6316

Centre for International Co-operation and Appropriate Technology (CICAT)
Stevinweg 1
NL-2628 CN Delft
Telephone +31 15 278 3612

Centre for Transportation Engineering
Stevinweg 1
NL-2628 CN Delft
Telephone +31 15 278 6634

Koiter Institute Delft (Institute for Engineering Mechanics)
Stevinweg 1
NL-2628 CN Delft
Telephone +31 15 278 5464

Institute for Wind Energy
Stevinweg 1
NL-2628 CN Delft
Telephone +31 15 278 5170

Interfaculty Reactor Institute
Mekelweg 15
NL-2629 JB Delft
Telephone +31 15 278 6712

OTB Research Institute for Housing, Urban and Mobility Studies
Jaffalaan 9
NL-2628 BX Delft
Telephone +31 15 278 3005

Open Building Working group (OBOM)
Berlageweg 1
NL-2628 CR Delft
Telephone +31 15 278 5400

Delft Institute for Microelectronics and Submicrontechnology (DIMES)
Feldmannweg 17
NL-2628 CT Delft
Telephone +31 15 278 3868

Interduct Delft University Clean Technology Institute
Rotterdamseweg 145
NL-2628 AL Delft
Telephone +31 15 278 3341

J.M. Burgerscentrum
Centre for Fluid Mechanics
Rotterdamseweg 145
NL-2628 AL Delft
Telephone +31 15 278 3216

Netherlands Schools for Advanced Studies in Construction
Stevinweg 1
NL-2628 CN Delft
Telephone +31 15 278 3332

Advanced School for Computing & Imaging
Mekelweg 4
NL-2628 CD Delft
Telephone +31 15 278 8032

TRAIL Research School
Kluyverweg 4
p.o. box 5017
NL- 2600 GA Delft
Telephone +31 15 278 6046

Central lLbrary

Delft University of Technology Library (DUTL) supplies information and provides services, particularly in the area of the technical sciences.
It comprises a central library and twelve sub-faculty libraries housed at the respective sub-faculties and institutes.
The DUTL is intended for students and staff at the Delft University of Technology.
However, as the task of the library is to provide scientific and technical information at a national level, its facilities are also available to the general public. As well as all areas of technology and natural sciences, the library also contains a general collection in the social sciences, economics etc.
This relates not only to books or periodicals, but also to standards, reports, reference works and congress proceedings.
Literature not in the collection or not on hand can be obtained through Delft University's Central Library from other libraries in the Netherlands or abroad.

For further information:
Delft University Central Library
Prometheusplein 1
P.O. BOX 98
NL-2600 MG Delft
Telephone +31 15 278 8022

Delft University Press

Delft University Press publishes scientific books and electronic magazines, such as the quarterly journal 'International Shipbuilding Progress' (LSP) (http://isp.tudelft.nl), European Journal of Transport and Infrastructure Research (http://ejtir.tudelft.nl) and the Journal of Design Research (http://jdr.tudelft.nl).
A catalogue is available free upon request.

Delft University Press
Prometheusplein 1
P.O. BOX 98
NL-2600 MG Delft
Telephone + 31 15 278 3254
Fax + 31 15 278 1661
dup@library.tudelft.nl

Information

General information:
Information office
P.O. BOX 5
NL- 2600 AA Delft
Telephone +31 15 278 5404
http://www.tudelft.nl

Information on facilities for foreign students:
Student Advisory Office
Julianalaan 134
NL- 2628 BL Delft
Telephone +31 15 278 6098

Liaison between business and research:
Liaison Office
Julianalaan 134
NL- 2628 BL Delft
Telephone +31 15 278 1500

Information on research fellowships:
Mrs. M.Y.M. Spiekerman
Julianalaan 134
NL- 2628 BL Delft
Telephone +31 15 278 3773

General information on university education in the Netherlands:
Min. of Education, Science & CultureCentral Information Dpt.
P.O. BOX 25000
NL-2700 LZ Zoetermeer
Telephone +31 79 353 1911

(Post Graduate) Courses

Delft TopTech
(vocational courses)
Mekelweg 2
P.O. BOX 612
nl–2600 AP Delft
Telephone +31 15 278 8019
Fax +31 15 278 1009
www.delft©toptech.nl

Institute for Biotechnology Studies Delft Leiden (BSDL)
Julianalaan 67
NL-2628 BC Delft
Telephone +31 15 278 6316

For information on courses in the Dutch language:
Language Laboratory
De Vries van Heystplantsoen 1
NL-2628 RZ Delft
Telephone +31 15 278 4124