Final Report
Perspectief Programme
Robust design of Cyber-Physical Systems (CPS)

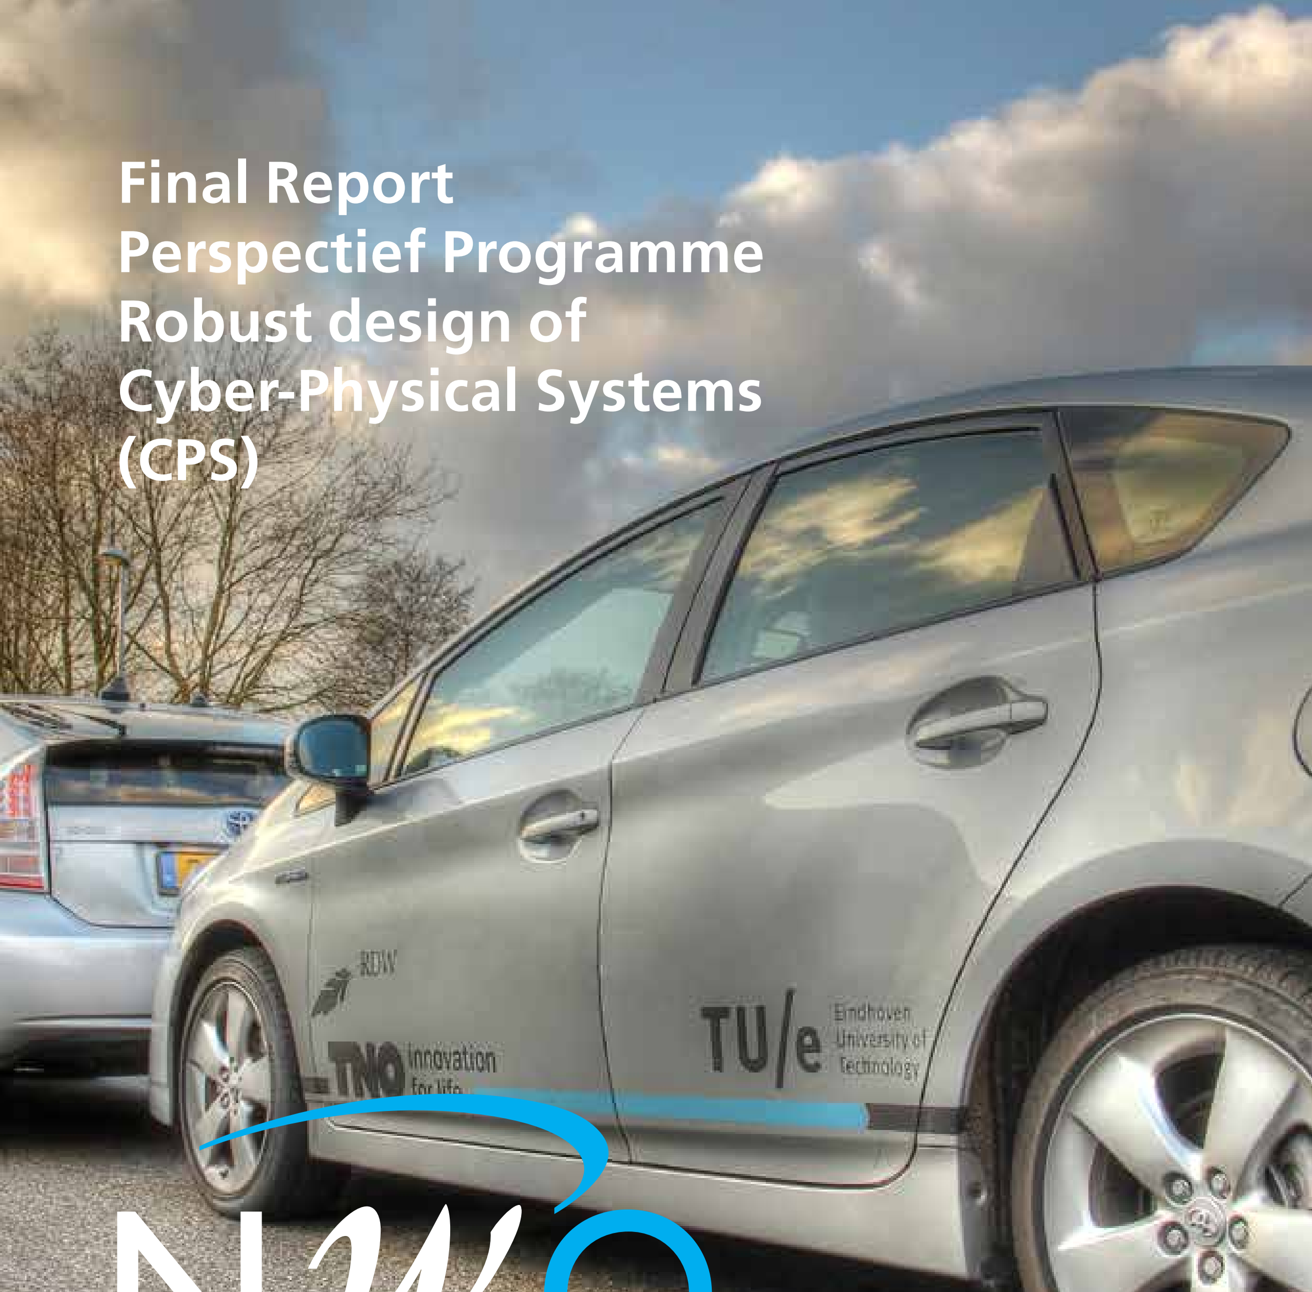

NWO

Applied and Engineering Sciences

Netherlands Organisation for Scientific Research

Final report
Perspectief Programme CPS

Robust design of Cyber-Physical Systems

Content

Combining computation, communication and control

Gerard Smit

A growing number of systems that execute physical functions like cooling, heating, printing or moving are controlled by computers. Examples include aeroplanes, printers or cooling systems in data centres. Designing these so-called cyber-physical systems – where embedded computers monitor, control and coordinate physical processes – is a growing challenge. The programme Robust design of Cyber-Physical Systems (CPS) has led to generic methods, tools and models to establish designs that enable fast, reliable, reproducible and robust control of such systems.

The design of next-generation, high-tech professional systems for application areas like medical imaging, lithography, smart electricity grids, intelligent transportation, electron microscopy and high-end printing requires a tight coordination between cyber aspects like computation, communication and control elements on the one hand, and physical processes such as heating, cooling, motion and vibrations on the other hand. Usually, these two parts of the system are designed separately and their joint performance is only tested for the first time when a prototype of the machine is built. This leads to unnecessary time delays in the development process.

As high-tech systems are becoming increasingly complex, workload variations and failures are getting more unpredictable, network connectivity is being stepped up, and time to markets are getting shorter, there is a need for new design principles that allow for simultaneous and interconnected design, development and testing of both the cyber and the physical part of the system. This is the precise focus of the CPS programme, programme leader Gerard Smit from the University of Twente explains: ‘In a way, the CPS programme is the successor of the NEtherlands STreaming (NEST) project that was granted in 2009 within Technology Foundations STW’s Open Technology Programme. In this project, we focused on products that operate on streams of data. This project was the largest ever granted in the Open Technology Programme and back then a consortium of companies struggling with the same sort of questions became heavily involved in the project.’

Taking control into account

During the realisation of the NEST project it became clear that the most prominent difficulties could only be solved if control is also taken into account at an early stage. 'So in 2010, we held a meeting in Utrecht to get new parties interested in this type of research. Together we wrote a new programme, which became CPS.' The programme was in every sense a forerunner at an international level, he says. 'There are some small comparable initiatives in countries like Germany, but none of them is even close to the size of the Dutch initiative, and industry is far less involved.'

The programme focused on a series of challenges, Smit explains. 'First, there is the complexity issue. There is an enormous growth in the number of sensors and actuators, each with their own data formats and time delays, which preferably have to be controlled in real time. That is especially important for applications such as robotic surgery or collision avoidance in automotive applications. Furthermore, you have to deal with uncertainties. What do you do when a sensor temporarily fails to provide you with information due to a hick up in the wireless communication network? How do you control a system when its behaviour, or that of its environment, is inherently unpredictable? How do you ensure that all parts of the system, whether it is a vibrating stainless steel part or a camera monitoring the position of a patient, communicate with each other with the required speed and type of information?'

The programme consisted of nine projects, each focused on a specific use case suggested by the partners involved. The projects were linked to five research lines, which addressed the generic scientific challenges, ranging from multi-domain modelling to energy-constrained sensing and actuation. Besides the standard user committee meetings on a project level, all researchers and users also met once a year at a programme level to discuss results and future directions. This integrated approach has led to success on multiple levels, states Smit. 'The science was of a very good quality: several of the PhDs gained their doctorates cum laude. Furthermore, the programme has led to two new Perspectief programmes being granted: the ZEROprogramme, in which we are looking for solutions to develop systems that can generate their own energy; and the Efficient Deep Learning programme, which is aimed at improving deep learning technology by reconsidering all design levels, from algorithms through to the hardware.'

Spot-on timing

Smit looks back on the CPS programme with nothing but pride. 'The timing of this programme was spot on: with the preceding NEXT project we had laid

the first foundations for a fruitful cooperation with industry, ensuring their commitment and contributions. All of the partners involved really shared the feeling that we were working on an urgent, joint problem. Within the individual projects we developed new models, tools and methods that are generically applicable and companies have shown a lot of interest in developing these further for implementation in their daily practice. In short, I think this programme has paved the way toward a more efficient approach for designing cyber-physical systems by pinpointing possible bottlenecks in a very early stage of development. This not only leads to higher quality systems, but also less time from design to market.'

10 years of Perspectief

You need to think big to realise breakthroughs in innovation. When scientists, companies and societal organisations work together then powerful forces can arise to tackle complex innovation bottlenecks.

That is what NWO's funding instrument Perspectief has been making possible for ten years. Dozens of research programmes to tackle major societal issues have started thanks to Perspectief. One such example is the programme 'Robust design of cyber-physical systems'.

The basis of each research programme is a consortium in which the entire chain needed to make the innovation possible participates. A consortium can consist of specialised start-ups with a brilliant idea, multinationals who can broadly introduce a new technology and organisations of end users.

In 2018, NWO celebrates Perspectief's tenth anniversary. For further information please see www.nwo.nl/10-jaar-perspectief. NWO realises Perspectief on behalf of the Ministry of Economic Affairs and Climate Policy.

Robust design of Cyber-Physical Systems

Budget

8 million euros, of which
3 million euros were contributed
by participating companies
and knowledge institutes

Size

nine projects
with 22 PhD students

Research lines

Hierarchical multi-domain modelling, analysis and synthesis, Stochastic analysis and synthesis, Control over communication networks, Distributed control, Energy-constrained sensing and actuation

Duration

2012-2018

Partners

Alliander NV, Altran BV, ASML Netherlands BV, BetterBe, Bosch Rexroth Electric Drives and Controls BV, Controllab Products BV, DEMCON Holding BV, Epilepsie-centrum Kempenhaeghe, ICT Automatisering Nederland BV, IMEC Nederland, Ipsum Energy, Malvern Panalytical, MathWorks, Medecs BV, Nederlandse apparatenfabriek 'NEDAP' NV, NXP Semiconductors, Océ Technologies BV, QBayLogic, Recore Systems BV, Sioux CCM, Target Holding BV, Technolution BV, Thales Group, Thermo Fisher Scientific (FEI), TNO-Automotive, TNO-ESI, Vanderlande Industries BV, Westnetz

Universities

Delft University of Technology
Eindhoven University of Technology
Radboud University
University of Amsterdam
University of Groningen
University of Twente

Project leaders

Prof. dr. ir. A.A. Basten
Eindhoven University of Technology
Dr. ir. J.P.M. Voeten
Eindhoven University of Technology
Dr. ir. S. Stuijk
Eindhoven University of Technology
Prof. dr. H. Corporaal
Eindhoven University of Technology
Prof. dr. C. de Persis
University of Groningen
Prof. dr. ir. W.P.M.H. Heemels
Eindhoven University of Technology
Prof. dr. ir. M.J.G. Bekooij
University of Twente
Dr. ir. J.F. Broenink
University of Twente
Prof. dr. ir. G.J.M. Smit em.
University of Twente
Prof. dr. J.J.M. Hooman
Radboud University

Projects Programme CPS

Robust design of Cyber-Physical Systems

Processing complex flows of objects

Prof. dr. ir. A.A. Basten, Dr. L. Somers **project** 12693

How can you process complex flows of objects in high-tech systems in such a way that you achieve the best product quality at the highest possible speed? Researchers from Eindhoven University of Technology, Delft University of Technology and Radboud University developed scheduling and control techniques that take cross-disciplinary constraints into account.

In many industrial machines, a flow of physical objects, such as wafers or sheets of paper, need to be managed under strict timing and quality constraints, with limited computational resources. To meet the demand for even faster, smaller, more energy-efficient and higher quality systems, this project developed model-driven design techniques to integrally design the scheduling of the objects, the distributed control, and the system's physical architecture.

Project leader Twan Basten from Eindhoven University of Technology explains what the project was about: 'We looked at the integrated scheduling and control of a particular machine of Océ. This printer is able to print the equivalent of two complete PhD theses in merely sixty seconds.' The project used cross-disciplinary reasoning strategies to optimise scheduling and system design, says Basten. 'Starting from the premise that you need to print the equivalent of 150 full-colour duplex A4s per minute without any prior knowledge about the print jobs, we tried to come up with a generic method to answer design questions like: What is the optimal geometry of the paper path?

Where and when should pages that have been printed on one side re-enter the incoming stream of sheets to be printed on the other side? How can unnecessary print head adjustments be avoided? And what heating strategy is best for the sheets of paper?'

Paper path design

At the Radboud University in Nijmegen, the researchers focused on machine-learning techniques to recognise patterns in print jobs. How do paper types and sizes vary? How often will there be jobs

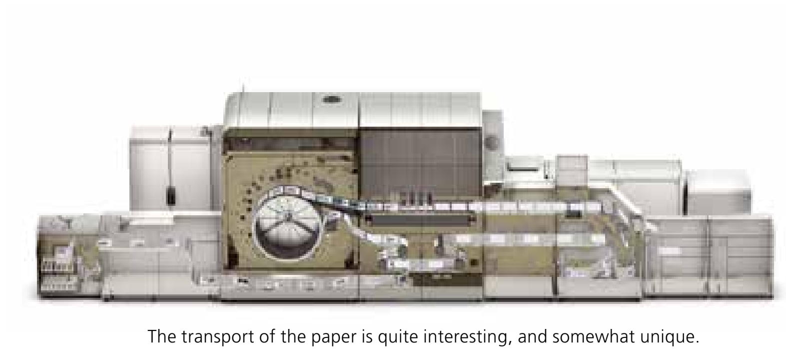

The transport of the paper is quite interesting, and somewhat unique.

The VarioPrint i300 is designed as a duplex machine, and as such has no advantage or penalty for running single sided or duplex print.

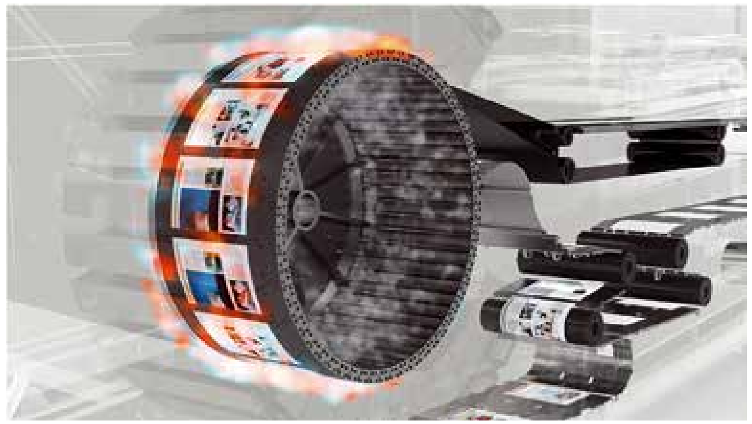

The drying drum uses infrared heating, air and a moderate drying temperature, to minimize paper deformation while still delivering a robust drying result.

containing book covers that are made of thicker, glossy paper that needs a different position of the print head? How many sheets of paper typically need to be printed on both sides? In Delft, the researchers delved into control problems: What is the optimal heat dosage to enable the ink to dry as fast as possible, without leading to paper damage? How to distribute the control over the machine? And in Eindhoven, the researchers focused on the sheet scheduling and paper path design.

Sheets of paper are fed into the machine, where they are printed and heated to dry in the printer's paper path. 'The nozzles release ink drops in the order of magnitude of picolitres. Since drops deform during their fall, the quality of the prints is determined by the height of the print head above the paper. Since adjusting the position of the print head takes costly seconds, scheduling sheets of different thickness is a challenge. Also size variations are challenging. We had to design a scheduling process that merges sheets that have been printed once already with empty sheets coming in, without leading to unnecessary time delays due to sheets that do not fit or due to print head adjustments.'

The project resulted in a constraint-based scheduling approach, says Basten. 'We started from a description of all constraints, such as sheet transport times, print head settings and the required minimal distance between two subsequent sheets. Then we developed algorithms for smart decision-making, which efficiently search through different scheduling options and based on the defined constraints quickly choose the optimal release time for the next sheet coming in.'

These algorithms are generically applicable, says Basten. 'For example, we are now also exploring their use in luggage handling and warehousing at Vanderlande Industries. Furthermore, we are working on integrating mechanical design parameters into the model so that mechanical bottlenecks can be identified in the design. This way we can make a shift from performance engineering in the final phase of development to model-driven performance design in an early stage. If the impact of different design choices can be explored at an early stage then lead times in product development will be shortened and a product can enter the market sooner.'

Researchers

Prof. dr. ir. A.A. Basten
Prof. dr. ir. B. De Schutter
Dr. ir. M.C.W. Geilen
Prof. dr. F.W. Vaandrager
Dr. ir. A.J.J. van den Boom
A. Firooznia MSc
A.R.R. Linard MSc
A. Rahimzadeh MSc
J.H.H. van Pinxten MSc

Users

ASML Netherlands BV
TNO Embedded Systems Innovation (TNO-ESI)
Océ Technologies BV
Vanderlande Industries BV

Collaboration crucial for long term innovation capacity

Lou Somers
software team leader at Océ Technologies BV

'We want to get more grip on the design of large complex systems such as our Varioprint i300. Traditionally, the paper path is mechanically designed following specifications about the reliability or cost, for example. Only when you start scheduling the paper do you find out that it would have been easier if the path had been ten centimetres longer, for example. But then it is too late.

This project resulted in useful insights about how design technologies work and how you can use them in the right way. One of my colleagues applied some of the research results in solutions that are now implemented in the next generation machine. In the future, we would also like to add machine-learning techniques to the scheduler. That way we hope to be able to build a system that optimises its own settings, based on its prior experiences.

These types of collaborations with universities and TNO are crucial for our innovation capacity in the longer term. In our daily business, we will apply and build further on the research results in order to realise product improvements at a higher level.'

Reasoning about design choices for complex high-tech systems

Prof. dr. ir. J.P.M. Voeten, Dr. ir. R. Schiffelers **project** 12694

High-tech systems like printers or wafer scanners need to produce products with increasing precision and speed. Manufacturers of such increasingly complex systems are looking for ways to reason about design choices that influence cost, accuracy and throughput at an early stage. In this project, researchers from Eindhoven University of Technology together with TNO-ESI and ASML developed methods and tooling to do just that.

Currently, different parts of complex high-tech systems are often designed and produced separately, and only when they are integrated can the system's performance as a whole be assessed. This may result in costly and time-consuming re-engineering efforts to meet the specifications set for the system. In this project, researchers developed a model-based methodology to explore the impact of design alternatives on the overall performance of the system at an early stage of development.

Juggling silicon wafers

The use case was provided by ASML and involved the wafer handling system of their lithography machines. In these machines, silicon wafers are transported through the machine to be illuminated by a powerful laser. After the illumination step the wafers are taken out of the machine to be developed and further processed, only to come back for a new illumination round to build the next layer of chip structures.

The logistics of this process is complex. The wafers should be moved and aligned in split seconds to maximize the number of wafers that can be illuminated per minute. But they also need to be positioned with nanometre accuracy to ensure that the chips produced will work as planned. A myriad of sensors, actuators, robots and software control loops is used to achieve this high precision, high speed, handling process. Due to this complexity, it is hard for a human to pinpoint which elements of this system lead to which part of the system's performance.

ASML wafer.

Three levels

'We studied this challenge at three levels,' explains project leader Jeroen Voeten. 'We went all the way up from controlling individual actions, such as "move the wafer from A to B", to complex logistics regarding the entire routing of wafers through the machines.' The project consisted of five different work packages, carried out partially at the university and partially at the premises of ASML. 'One of these work packages, for example, focused on mechatronics, to pinpoint which sensors in the system are critical and need to be of the highest quality, and which can be less costly. We also studied methods for design-space exploration: what is the impact on throughput if you add an extra wafer-handling robot? And at the logistics level, we could in the end take a set of requirements and automatically deduce the required overall controller to achieve optimal throughput.'

The project resulted in methods and tooling that can be used to calculate the optimal throughput given a certain combination of hardware components and control software, to pinpoint which components are critical for the overall system performance, and to get an idea of the expected system performance based on true components.

Tight links with TKI projects

Voeten looks back at the project with pride. ‘Our tools are currently used by ASML to design their next-generation machines. And since the tooling we developed is generic, it can also be used by Océ, for example, to optimise the paper paths in their printers.’ The secret to this success was the fact that this project was closely related to two other joint projects with ASML, emphasises Voeten. ‘In the TKI projects Concerto and Carm2G, we could employ researchers from TNO and scientific programmers and industrial architects from ASML who could translate our basic research findings into something that is applicable in an industrial setting.’ This also works the other way around, Voeten says. ‘Those two projects helped us as scientists to deduce the scientific questions that lie beneath the practical challenges.’

Beyond happy flow

The cooperation between the different parties involved will be continued, says Voeten. ‘Until now we focused on happy flow behaviour, studying the behaviour when systems are in production mode. But this is just a single scenario of the behaviour a machine can be engaged in. As a next step we will broaden our scope to take other scenarios into account, such as maintenance, calibration and error handling, and especially the interaction between these. If we can integrally specify and analyse these scenarios and their interaction, we will not only be able to optimise overall system performance, but also to bridge the gap between multidisciplinary specification and monodisciplinary implementation. This will have a major impact on engineering efficiency and effectiveness, and help our high-tech industries to meet their demanding time-to-market and time-to-quality constraints.’

Researchers

Prof. dr. ir. J.P.M. Voeten
Prof. dr. ir. A.A. Basten
Prof. dr. H. Corporaal
Dr. ir. M.C.W. Geilen
Dr. ir. S. Stuijk
J.P. Nogueira Bastos MSc
L.J. van der Sanden MSc
Prof. dr. ir. W.P.M.H. Heemels
Dr. ir. M.A. Reniers
Dr. ir. T.A.E. Oomen
J.C.D. van Zundert MSc
Prof. dr. M.G.J. van den Brand
Dr. A. Serebrenik
Dr. ir. R.R.H. Schiffelers
J.G.M. Mengerink MSc

Users

ASML
TNO-ESI
ICT Automatisering
Océ Technologies
Sioux CCM
Technolution

Exploring feasibility of specifications

Ramon Schiffelers
senior software architect at ASML

‘At ASML, we wanted help in designing the logistics controllers for our wafer handling subsystems: What parts have to move at what time, via which route and in which order? As our systems become more complex it is increasingly difficult to gain insight into which design choices will lead to the optimal machine output at affordable costs.

The project has resulted in useful tooling to design our machines. We can now model new machine types and get an idea of the feasibility of a certain envisioned throughput. Not only we, but also our suppliers can use the tool in an early stage of development to see which parts of the system are critical. Which parts are worthwhile improving and for which parts won’t it make any difference? At ASML we are going to extend the tooling to include product line engineering as well. We applaud the initiative taken by the university to incorporate the methods and tooling devised into their education. We see a lot of added value in students who are trained to analyse systems at a more abstract level.

The existing collaboration with TNO-ESI and Eindhoven University of Technology works really well, and is something we will continue in the future. After all these years we understand each other’s needs and value each other’s expertise, and our cooperation has proven to lead to good science as wel as useful solutions for our daily practice.’

Energy-efficient sensor networks

Prof. dr. H. Corporaal, M. Wijtvliet Msc., Dr. ir. G. Rauwerda **project** 12695

A growing number of sensors are being developed for medical purposes that can measure body functions, either on or inside the body. In these systems, battery replacement is often intrusive. Therefore, to prolong battery life, processing and communication processes should be aware of the available energy and use it efficiently. In this project, researchers from Eindhoven University of Technology, Leiden University and the University of Twente developed generic methods, tools and protocols for an energy-efficient, multi-sensor, processing platform.

The project focused on EEG monitoring, which exploits subcutaneous implanted sensors to predict epileptic seizures, explains project leader Henk Corporaal from Eindhoven University of Technology. The final envisioned system should consist of sensors implanted under the skull combined with a baseball cap or hair band that both provides energy to the subcutaneous sensors and enables read out and processing of the high-quality EEG signals captured by them. An alternative solution is to use a wearable headset with several electrodes. This method does not require implants, but the signal quality is much lower and so far, more signal processing is required.

The project addressed several challenges that have to be overcome to realise such a system, Corporaal and PhD student Mark Wijtvliet explain. 'Our colleagues at the University of Twente focused on body sensor networks and how these communicate with the outside world. For example, they characterised how radio waves travel through people, since that is how wireless implanted chips communicate. They also studied how the signal strength depends on the movement of specific body parts and then deduced when and where to make measurements and to communicate the data gathered.' Their efforts resulted in new energy-efficient and reliable communication protocols that can deal with data loss due to wireless link deterioration. At Leiden University, the researchers tried to discover how to schedule multiple tasks in a throughput-constrained streaming application with limited resources, memory and number of processors. Their research resulted in a new scheduling approach and in algorithms that utilise processing resources more efficiently, while still guaranteeing throughput with due consideration for latency constraints.

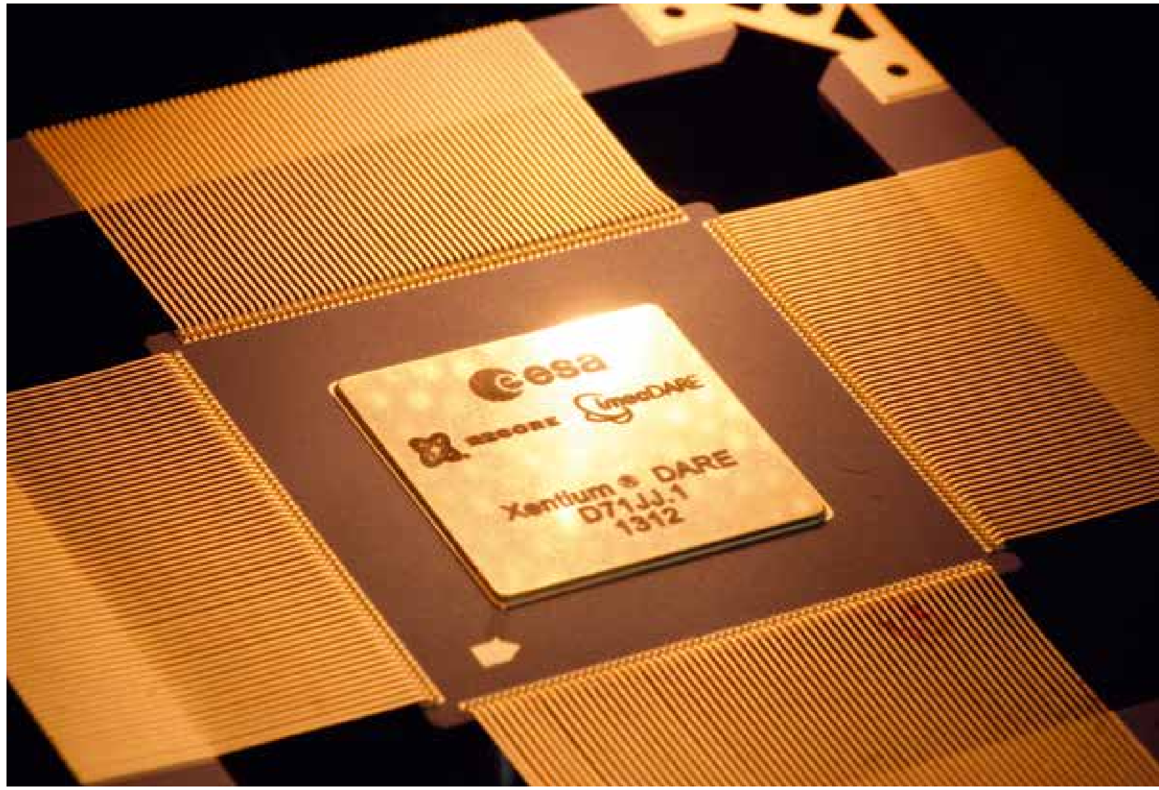

Radiation-hardened Digital Signal Processor prototype
for on-board payload data processing © Technolution BV.

Neural networks on mobile platforms

At Eindhoven University of Technology, two PhD students looked into the data processing part. How do you automatically conclude from an EEG if a seizure is indeed developing, and how can you implement the neural networks that can do this on a low-power mobile platform? 'The main bottleneck for performance and energy efficiency is the external memory traffic,' explains Corporaal. 'We transformed the required scheduling of processing operations in such a way that most of these can be performed by the internal memory instead of the external one. This leads to a reduction in energy consumption and improved performance,' says Wijtvliet.

Furthermore, the Eindhoven researchers designed a new chip architecture called Blocks, which reduces the amount of communication needed. ‘We devised a chip containing multiple flexible processing elements that can be configured by software to meet the demands the application sets on it,’ explains Wijtvliet. ‘That way the software can define the best processor lay out for the job at hand, leading to shorter processing times and energy savings.’

The first prototype of the chip has been built. And only minutes before the interview, Wijtvliet saw proof that it actually works. ‘We still need to work on the compiler though,’ he emphasises. ‘But for now, it looks as though we have developed a generic, flexible architecture for signal processing at low power, which is over 300 times more energy efficient than comparable processors. The only drawback is that it is slightly larger because it is reconfigurable and therefore contains more components.’ However, as the chip allows for parallel processing, it is just as fast.

Even though some of the PhD students have not finished their work yet, a follow-up project has already started. ‘In the BrainWave project that we are running together with Kempenhaege and the Donders Institute, we are now including patients and experts in EEG data classification. In this project we are not only looking at epilepsy, but also at Parkinson’s disease,’ concludes Corporaal.

Researchers

Prof. dr. H. Corporaal
Dr. ir. N. Meratnia
Dr. ir. T.P. Stefanov
Prof. ir. ing. Paul Havinga
Ir. Jos Huisken
S. Niknam MSc
V.R. Karuppiah Ramachandran MSc
L.J.W. Waeijen MSc
M. Wijtvliet MSc

Users

Radboud University Department of Biophysics
Recore Systems BV
Medecs BV
CED/Diagnostics
IMEC Nederland
NXP Semiconductors
Epilepsiecentrum Kempenhaeghe
TMSi
Donders Institute

Low-power processors for on-board payload processing in space

Gerard Rauwerda
Business Developer at Technolution, previously CTO at Recore Systems

‘At the time this project started, I worked at Recore Systems where we developed reprogrammable digital signal processors. This project had the same aim, but explored a different technology to achieve it. Since this other technology probably could enable other functionalities, we were very interested.

Unfortunately, a couple of months ago Recore Systems was declared bankrupt. Technolution has taken over some of Recore’s activities, especially those aimed at digital signal processing for space applications. Typical space exploration missions collect large amounts of data. Since the communication bandwidths are not high enough to transmit all of this data back to earth, part of the data processing is done in space. This must be done in an energy-efficient manner. This is not due to a lack of energy available but because chips that use a lot of energy heat up and cooling is a big problem in space, since there is no air.

This project has resulted in an interesting new architecture for low-power, reconfigurable processors. These results will need some extra work before they are really ready for implementation though, especially on the compiler and software side. However, in my new role at Technolution, I will maintain close connections with universities to keep track of the latest technological developments and to scout innovations that are relevant for our daily business.’

Saving energy in datacentres

Prof. dr. C. de Persis, T. van Damme Msc., Dr. ir. G. Meester **project** 12696

Datacentres, the gigantic buildings where your Google searches and Facebook friendship requests are processed, currently use up more than three percent of all of the electricity we generate worldwide. If we don't act now, this figure will only go up in the coming years. In this project, researchers from the universities of Groningen, Twente and Amsterdam explored ways to drastically reduce the power consumption of future datacentres.

The project investigated possible energy saving methods on three different levels, PhD student Tobias van Damme from the University of Groningen explains. 'My colleagues at the University of Amsterdam dived into possible improvements at the level of the central processing units (CPU) within the servers. In datacentres, multiple CPUs can be working on a single task simultaneously. By dividing the workload in a smarter way, energy can be saved.' The challenge there is to guarantee all tasks are executed in the most energy-efficient manner, without leading to loss of information or time delays. The researchers developed and tested new scheduling algorithms, and studied a datacentre model for power-aware scheduling of applications with timing constraints.

Virtual testbed

At the University of Twente, a virtual testbed was developed that can be used to simulate the effect of different energy-saving solutions, Van Damme says next. 'For example, they showed potential energy savings and the existence of trade-offs that occur in a scenario where servers are switched on and off depending on the demand for services. Obviously, by turning off a machine, you save the power that is used by a system in a standby state. But this is only worthwhile if you turn them off for a long enough period of time, since rebooting consumes a lot of energy and time. On top of that, there is always a risk that such a server won't reboot because of some error, leading to unwanted delays in performing the demanded tasks. A more feasible solution analysed is the possibility of putting servers into sleep states reducing booting time and unwanted booting errors.'

Datacentre.

In their simulations, the researchers calibrated their models with data from realistic workload traces from project partner BetterBe. 'They provide cloud-based services for the global mobility and car leasing market, which enable clients to combine desired options and immediately see what the final car would cost them on a monthly basis. These are all relatively short, small tasks, which require a lot of communication between different databases and a lot of computation performance.'

Take the heat off

Van Damme himself conducted a theoretical study of the thermodynamics within the datacentre. 'A datacentre consists of a series of racks containing the servers. These racks are set up in aisles. Central cooling units blow in cold air to cool down the servers on the front side of the racks. The heat is dispersed on the rear side. There is no possibility to cool down specific servers, so the amount of cooling is adjusted to the hottest location, which should be kept below a certain threshold.' Van Damme worked on a model which describes the thermodynamics of an entire datacentre. He used air circulation information provided by IBM Zurich to characterize the air flows inside the datacentre. 'From this model we designed and analysed controllers that divide the workload over servers in such a way that the temperature distribution becomes more gradual, leading to less high peaks in temperature, and thus to lower cooling demands.'

Van Damme's thermal-aware control has been implemented and successfully tested in the simulation tool at the University of Twente. 'Now it would be nice if we could also incorporate the CPU models from Amsterdam in the simulation tool, and run all of our solutions in a large-scale simulation of an actual datacentre. If our ideas lead to an energy reduction of only a few percent, we will have earned back the personnel costs of this project within one year of operation.'

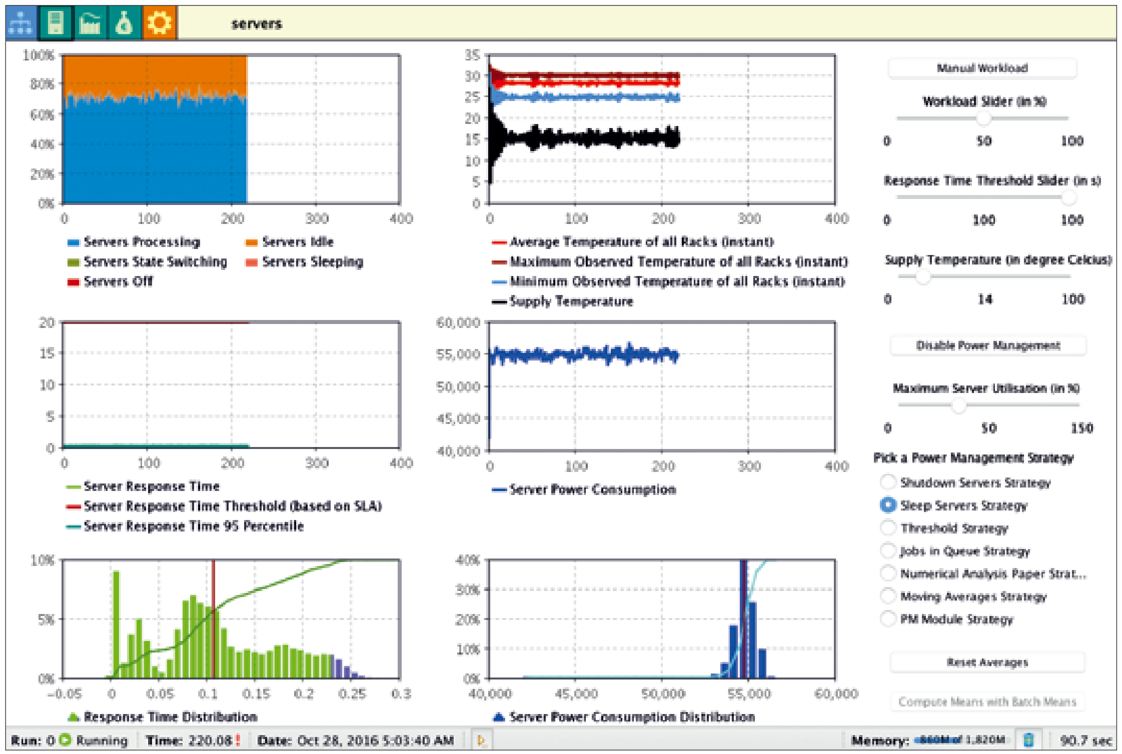

Researchers

Prof. dr. C. de Persis
Prof. dr. P. Tesi
Prof. dr. A.D. Pimentel
Prof. dr. ir. B.R.H.M. Haverkort
J. Xiao MSc
B.F. Postema MSc
T. van Damme MSc

Users

BetterBe BV
Target Holding BV

Searching for significant solutions

Geatse Meester
CEO BetterBe BV

'When I was approached to cooperate on this research project, I immediately said yes. As a company that offers solutions to transform the global mobility and the automotive leasing market, we own three identical datacentres. In the past five years, we have managed to drastically reduce their energy consumption, for example by implementing natural cooling processes. The main eye-opener of this project was the insight about which measures will lead to significant improvements and which will only yield marginal savings. For us the most practical solution was the idea of managing the uptime of our servers in times of low demands. When should we put them in hibernate modus and when can we best turn them off completely? Of course we will not implement this until we are 100 percent sure that we can guarantee all client requests will be processed on time without any hiccups.

During the project, we got more and more involved. Our datacentres are loaded with sensors, and the information they supply turned out to be very valuable for the research. Students have regularly worked with our employees and we have given them access to our anonymised client data to get an idea of actual workloads. The fact that we were willing to invest the valuable time of our own people in this, reflects the relevance of this research project for our business.'

Using noisy data to realise high-performance control

Prof. dr. ir. W.P.M.H. Heemels, Dr. ir. B. Janssen **project** 12697

In multiple high-tech systems and applications there is a drive towards accurate and fast control based on large amounts of data of diverse quality. At Eindhoven University of Technology, researchers developed communication and control algorithms to achieve real-time and high-performance control based on bundles of data, which are unpredictable in terms of quality and processing time.

Thermo Fischer Scientific provided one of the use cases for this project explains project leader Maurice Heemels. This company develops electron microscopes that enable imaging at an atomic scale. Because of the high resolution, minor movements caused by thermal vibrations or drift of the sample with respect to the stage lead to blurred images. If you can wait for a couple of hours after you have loaded your sample until the system has stabilised then there is no problem. 'But these electron microscopes are going to be increasingly used for faster measurement procedures, for instance, to quickly check the quality of very small structures in semiconductors. In that case, you want to be able to control the microscope in such a way that it can automatically keep its focus, despite the presence of vibrations and drift.'

A major challenge in this specific case is the microscope acting as the sensor to determine the image quality. In fact, the image blurring microscopic movements of the objects you need to focus on are only visible with the microscope. Therefore, the information gathered by the "microscope" sensor needs to be fed into a real-time control loop. You need good images for optimum control but at the same time you need optimum control to obtain good images.

Time versus accuracy

An important design decision is to weigh up processing time to obtain information from the data against the accuracy of the information obtained and also to deal with the chance that there will be no information at all, explains Heemels. 'Drift information

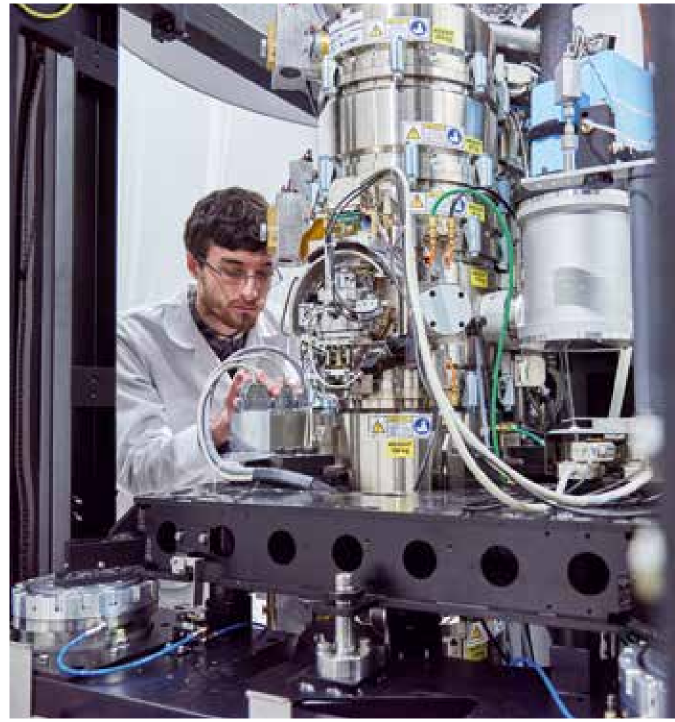

Electron microscope.

from an image produced by an electron microscope is extracted by analysing the entire image. This leads to a long processing time and a high probability of relevant information with high accuracy. Alternatively, you could analyse just a few lines or pixels of the image, leading to short processing time, but less accuracy. The question is how long you should wait for the data-processing algorithm before you take actions to counteract the drift by moving the stage.'

The researchers started by developing an integrated modelling framework that combined the motion dynamics of the sample and the stage with models for data-processing methods including the trade-off decision between the processing time, the chance of information being gathered, and the accuracy of that information. 'This integrated modelling framework combined with the actual data and models provided by Thermo Fisher Scientific were instrumental in determining on-line which control algorithm will work best at any given moment during the operation. This leads to an adaptive control algorithm that will sometimes decide to adjust the position of the stage based on one image line, and sometimes waits until the entire image is analysed.'

Resource usage

The researchers also looked into optimising the use of processing and communication resources. Multi-core processors enable multiple cores to be allocated to the image processing for pipelined processing of the image data. This allows a trade off between quality of control versus processor usage. To make effective use of the communication bandwidth in a control loop, the researchers looked into shaping the data communicated and improving the real-time behaviour of the communication.

'Interestingly, our solutions are not limited to electron microscopes,' emphasises Heemels. 'They are also highly relevant for application areas such as vision-based robotics or automated driving, where data from many different sources such as sensors, traffic lights or cameras is used to steer actions. Together with Technolution, one of the other project partners, we also looked into the automated driving case. And in our robotics lab, we have built our models into a demonstrator for applications in vision-based robotics. This shows the broad impact our results could have for our society.'

Thermo Fisher Scientific Lab.

Researchers

Prof. dr. ir. W.P.M.H. Heemels
Dr. D.J.G.T. Antunes
Prof. dr. ir. A.A. Basten
Dr. S. Stuijk
Prof. dr. ir. J.J. Lukkien
Dr. ir. .J.L. Cuijpers
R.A. Medina Sánchez MSc
Dr. ir. E.P. van Horssen
J. Cao MSc

Users

Technolution BV
Thermo Fischer Scientific
Bosch Rexroth Electric
Drives and Controls BV
MathWorks

Improving pictures taken in the dark

Bart Janssen
Staff Scientist Thermo Fisher Scientific

'Our microscopes are able to make films of processes that take place on an atomic scale. This requires short sampling times and active compensation for disturbances such as sample drift and vibrations. That is extremely challenging because many of the disturbances can only be observed by processing data provided by the microscope's image sensor, which typically produces large amounts of noisy data. With an electron microscope, it is often like trying to take a picture in the dark. You need a long exposure time to be able to see anything at all. But if you take too much time, your picture will be blurred by motion, part of which cannot be corrected for in retrospect. We therefore need to design algorithms that extract information from the images collected by the microscope and subsequently use this information to optimise the data collection process. This is particularly challenging because it needs to be done during image acquisition.

To make sure the PhDs working on this project obtained a good idea of this and other image optimisation challenges, we subjected all three of them to a week-long training, turning them into certified electron microscope operators. I think this investment at the start of the project was worth it. We are very happy with the results that have been achieved. The models that were developed turned out to be very useful additions to our simulations. During the project we even invested in an extra PDEng for two years to translate some of the models into practical solutions for our future machines. He did so well that we recently hired him.'

Improving safety-critical communication between cars

Prof. dr. ir. M.J.G. Bekooij, Prof. dr. ir. W.P.M.H. Heemels, Dr. ir. J. Ploeg **project** 12698

Wireless car-to-car communication systems combined with automated brake and throttle can enable platoons of vehicles to drive at short distances from each other. This saves fuel, helps to prevent phantom traffic jams and enlarges the capacity of the road. But how do you guarantee robust automated driving in platoons when the communication bandwidth is limited and can be unreliable due to packet losses and delays? In this project, researchers from Eindhoven University of Technology and the University of Twente developed new tools and methods to enable reliable and resource-aware car-to-car communication.

To create a platoon of automated vehicles, the information the individual cars receive via the wireless communication system needs to be combined with information gathered by other sources such as radar and optical sensors. Timing is crucial here: if the first vehicle in line communicates that it is going to brake, this information should result in the vehicles driving behind it braking at the same time. But if all vehicles are continuously communicating with each other, the communication network will soon get overloaded and essential information might not get through in time.

To communicate or not to communicate

At Eindhoven University of Technology, PhD student Victor Dolk studied the possibility of using event-based instead of time-based communication, explains his daily supervisor Maurice Heemels. ‘Time-based communication means that every car sends information at fixed equidistant times, even if nothing important happens. This is current practice and works well at the moment. However, if cooperative driving is implemented on a large scale this could lead to an overload of the communication networks. Victor Dolk developed a method that only lets the cars communicate when something hanges, for example when a car brakes or accelerates. Essentially, the relevance of the information present in measurements is assessed and this determines whether it is necessary to communicate this information over the network or not to realise good driving behaviour:

Wireless car-to-car communication system © TNO.

To communicate or not to communicate, that is the question!’ Dolk not only proposed this method and developed the necessary mathematical algorithms to achieve it, but also implemented his solution on a platoon of cooperative and automated cars at TNO in Helmond. ‘In those experiments, the first car was in the lead and communicated relevant changes in acceleration to the rest of the platoon. All of the following cars braked and accelerated automatically and (almost) simultaneously, leading to string stable driving behaviour,’ Heemels says. ‘Furthermore, to illustrate that event-based communication is not only relevant for automated driving applications but for any kind of system of system coordinating through wireless communication networks, we also implemented the algorithms successfully on the robots of the world champion TU/e Techunited soccer team.’

Analysing latencies

At the University of Twente, two PhD students delved into the specifications of the modem used for the car-to-car communication, explains project leader Marco Bekooij. 'The chips inside such a modem should be able to perform fast calculations. They need to consider not just the sampling rate, but also the maximum amount of time it may take before a result needed to control the vehicle's behaviour is required.' The PhD students developed analysis tools to assess if the modem is fast enough and to determine if guarantees can be given about its performance. Not only do all of the separate components on the chip need to be fast enough, but also the entire car-to-car communication process needs to be reliable, says Beloin. 'Imagine that a data packet gets lost along the way. Usually there is a handshake protocol in place that detects something went wrong and the packet is resent. But then a longer time has passed between one car sending the information and the other one actually receiving it. My colleagues in Eindhoven looked into what the consequences are if something like this happens: under what conditions in terms of communication rates, latencies and packet loss behaviour will the control system remain safe and stable, and under which conditions will safety be at stake? At our end, we looked into possibilities to control this latency.'

Fundamental theory

PhD student Philip Kurten therefore developed an end-to-end latency analysis theory, says Bekooij. 'The beauty of this theory is that it is not only applicable for discrete signals, like the ones processed by a chip, but also for hybrid signals that combine discrete and continuous parameters. This is very relevant for cyber-physical systems, which often combine both. The theory predicts when models are trustworthy and are in fact saying something about the actual performance of a device.'

Guus Kuiper, the other PhD student in Twente, developed a real-time data flow analysis tool. 'His research resulted in the insight that it is fundamentally impossible to obtain an exact analysis,' Bekooij explains, 'since there is a trade-off between precision and calculability.' Both Heemels and Bekooij state that this project has been very successful. 'It led to a cum laude PhD for Victor Dolk (which is extremely rare at Eindhoven University of Technology), important new fundamental insights about how to organise communication between safety-critical systems, and both NXP and TNO Automotive were very happy with the results. NXP is implementing some of the tooling developed in Twente in their radar systems for cars and the event-based algorithms from Eindhoven will be further explored in TNO's cooperative driving vehicles.'

Researchers

Prof. dr. ir. M.J.G. Bekooij
Prof. dr. ir. W.P.M.H. Heemels
Dr. ir. V.S. Dolk
G. Kuiper MSc
P.S. Kurtin MSc

Users

TNO Automotive
NXP Research & Development
Technolution BV

Prevent congestion of communication channels

Jeroen Ploeg
Lead Cooperative Driving at 2getthere in Utrecht, formerly Principal Scientist at TNO Automotive

'At the time this project started, I was working on integrated vehicle safety at TNO Automotive. Just like in most current automated driving communication systems, we used time-based communication: our cars exchanged information every 40 milliseconds, even if nothing important happens. One of the nice results of this project was that it showed the feasibility of event-based communication, which reduces the communication load drastically without decreasing the performance of the platoon in terms of string stability and passenger comfort. An important issue has not been fully addressed yet: what happens if the communicated signal does not reach the intended receiver? In a time-based situation, you will know that immediately. With event-based communication, that is harder. I can imagine that it will always be necessary to exchange some sort of heartbeat, to let each other know the connection is still working properly. This forms an interesting question for future research that could further underline the relevance of event-based communication. During the project I moved to the company 2getthere, where we develop automated transport systems. Since available communication channels have sufficient bandwidth for the number of cooperative vehicles driving around now, we are still implementing time-based communication protocols. But for the future, when the number of communicating cars is growing tremendously, event-based communication will play an important role. Therefore, it is good that this solution has been suggested and implemented in a proof-of-concept study. This idea should certainly be developed further.'

Model-driven software design for complex cyber-physical systems

Z. Lu MSc, Dr. ir. J. Kuper **project** 12699

How can you develop predictable and reliable software for complex cyber-physical systems in a systematic and semi-automated way? At the University of Twente, researchers Jan Broenink and Zhou Lu used a model-driven approach to develop a robust methodology for designing simulation software and control software.

In the design of robotic systems containing a software controller and a plant, the interfaces between the cyber part and the physical part are often not described precisely enough. This can cause serious problems during system integration, especially when dealing with provisions to ensure the safety of the cyber-physical system. Furthermore, a design approach that incorporates early testing and automatic code generation can lead to better products being developed in shorter timespans.

'In this project, we focused on combined modelling and coordinated simulation,' explains PhD student Zhou Lu. 'We combine models from different domains to describe the system, and verify the system as a whole with co-simulation tooling.' The problem here is that many tools or models can be used for design and simulation, says Lu. 'Our challenge was to develop a system that automatically combines these tools into one single platform.' Since every application has its own syntax and semantics, the researchers developed methods to transform the models in such a way that they could understand and complement each other.

Master-slave architecture

'We decided to use the Functional Mock-up Interface (FMI) standard to interconnect these subsystem models, since currently, over 100 tools claim to support this standard.' A model implementing such an FMI interface is denoted as an FMU, says Lu. 'We connected several of these FMUs in a so-called master-slave architecture, to end up with a co-simulation tool.' The slaves are made up of the FMUs that contain the individual simulation software. 'Furthermore, we

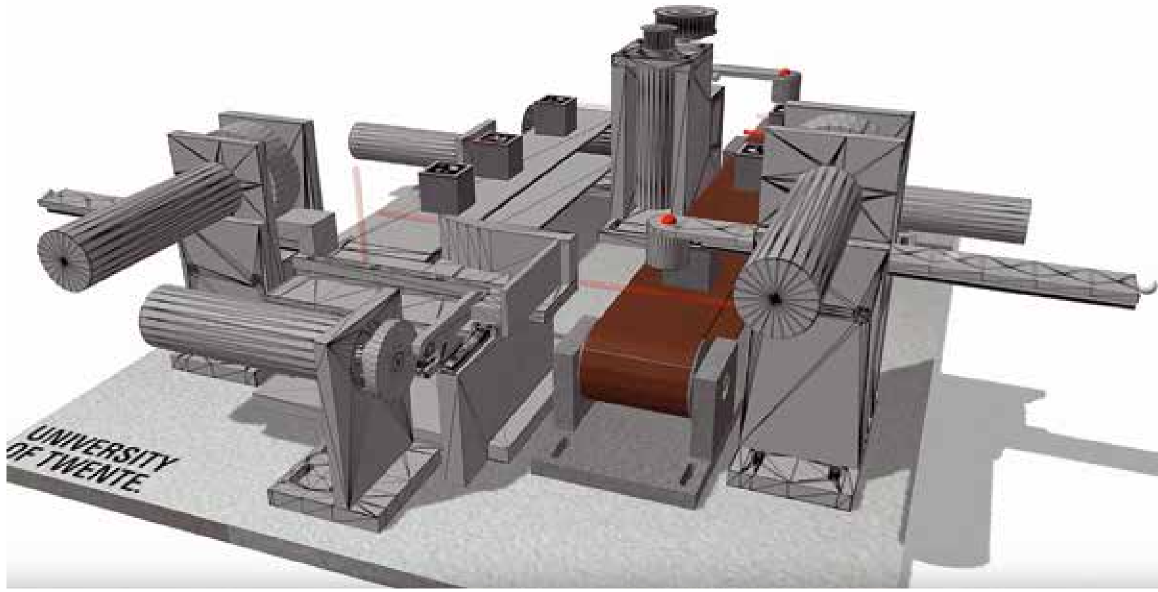

One of the demonstration models that was built:
a 3D simulation of an automated package sorter.

developed a master algorithm that determines the order of execution and synchronises the different datasets used by the different simulation tools.'

'One of the bottlenecks in co-simulation tooling is that it is often very slow, due to the amount of communication needed between the different simulation models. We developed ways to significantly speed up this process, by devising efficient solutions for data exchange.' The necessary software code to operate all of the simulations together is directly generated from the overarching model. The nice thing about Lu's architecture is that all sorts of tools can be integrated in it, such as models for communication channels or 3D-animation tools for robotics.

Demonstrators

Lu shows one of the demonstrators built within the scope of the project. 'This is a 3D animation and the ancillary performance simulation of an automated production facility that MSc student Wilbert van de Ridder developed based on our models and tools.' The plant is a package sorter, which transports and scans packages of different shapes and sizes, and throws out specific packages that are marked for this purpose.

‘We developed a proof-of-principle for model-based, co-modelling and co-simulation,’ Lu concludes. ‘Our solutions are already available for use in a research and education environment. Further utilisation has been boosted by start-up QbayLogic. This company is incubated at DEMCON, which took part in our user committee.’

A second demonstrator that was built:
a tunneling ball device.

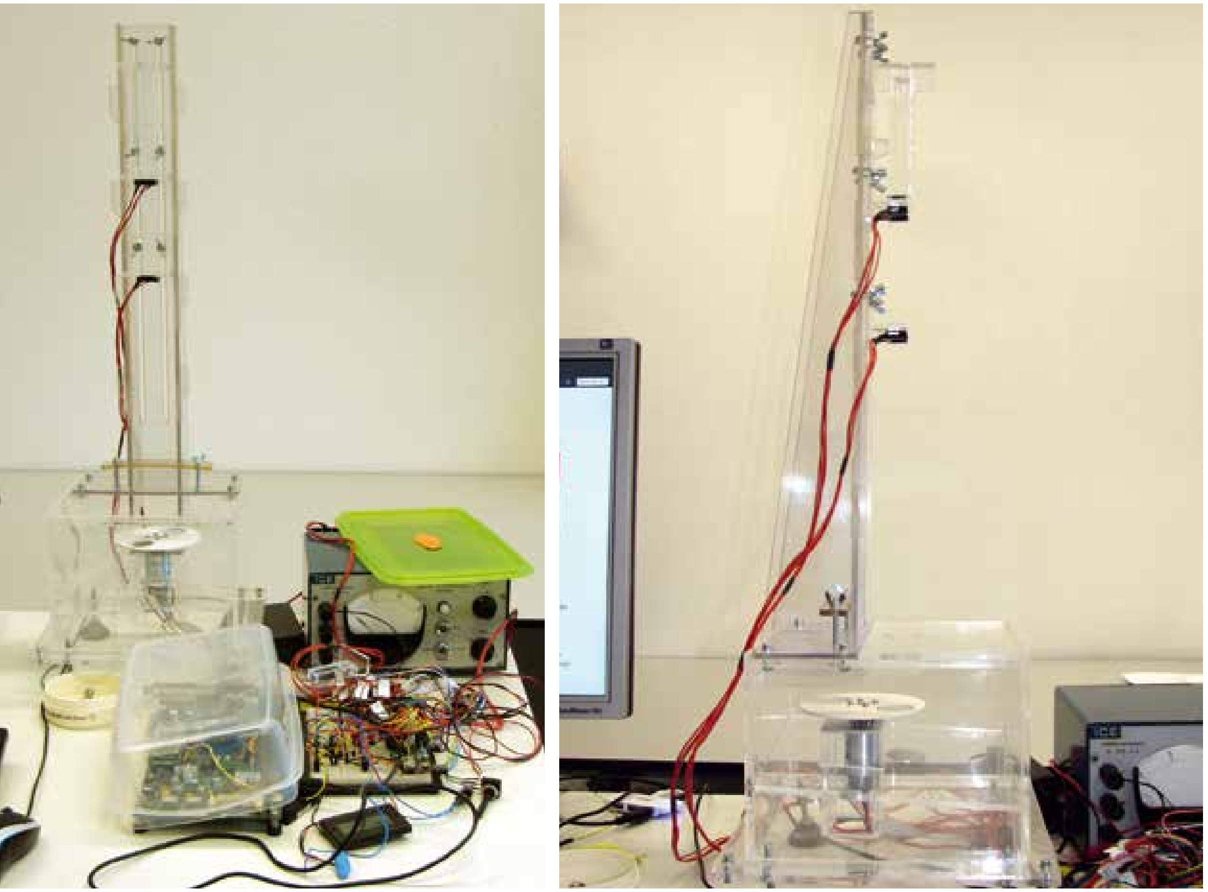

Researchers

Dr. ir. J.F. Broenink
Z. Lu MSc

Users

Altran BV
DEMCON Holding BV
Nederlandse apparatenfabriek 'NEDAP' NV
QBayLogic
Thales Group

Demonstrating the power of high level programming

Jan Kuper
Executive Director QBayLogic

‘As a start-up of the University of Twente, we offer tools to program so-called field-programmable gate arrays (FPGA): flexible computer hardware that can be programmed to act as a switch or any kind of logical component depending on the task at hand. Because the computing is done in the hardware itself, this is faster, cheaper and costs less energy than in traditional architectures. We program the hardware using CλaSH, a language that borrows both its syntax and semantics from the functional programming language Haskell.

In this project, this language has been developed further and some demonstrators have been built to illustrate its power for real-time control of cyber physical systems. One of the demonstrators built by student Peter Lebbing was a tunnelling ball device. This is a system that needs fast control to make sure a randomly dropped ball exactly falls through a hole in a rotating disc. With this system we can convince people working on time-critical systems, such as medical devices or laser communication between satellites, that CλaSH is the way forward if you want to use FPGAs.’

Sharing energy with your neighbours

Prof. dr. ir. G.J.M. Smit, Ir. P. van der Sluijs **project** 12700

When enough electricity generators like solar panels, biomass installations and wind turbines are installed and storage facilities are in place, it is possible to create a self-supplying neighbourhood in a so-called energy autonomous, smart micro-grid. In this joint project of the universities of Twente and Groningen, researchers developed models and software to efficiently balance energy supply and demand in such micro-grids.

Since the energy supply from renewable sources like wind turbines and solar panels fluctuates and is inherently uncontrollable, we are faced with severe challenges when controlling energy supply to match local demand. Under these circumstances, micro-grids can be an interesting alternative for large electricity networks. 'Not only is it easier to match supply and demand in smaller networks, but if you use the electricity that you generated yourself then you do not have to transport it and that prevents energy losses and costly grid investments that are only needed occasionally along the way,' says project leader Gerard Smit.

One of the questions addressed in the project was how to keep the 50 Hz electricity supply stable under fluctuating conditions. 'In a traditional electricity plant, the 50 Hz frequency is kept constant by making use of the inertia of the huge generators, which filter out a lot of the smaller fluctuations. But in a micro grid, you need electronics to control this frequency. Within the project, we developed models and control algorithms to tackle this problem,' says Smit.

Completely self-sufficient

Other modelling work was done on configurations for micro-grids. 'There, we not only looked at the consumption of electricity, but also of heat,' explains Smit. 'We have proposed a concept consisting of a quadrant of 16 houses, where every house generates electricity with solar panels. They share a battery to store the surplus of electricity generated on sunny days, and a combined heat and power unit that generates heat to warm tap water and the houses.

Windmills on sea.

We demonstrated that with this concept these 16 houses can be fully energy autonomous throughout the year and that they only need a small 2 kilowatt hour battery per house to achieve this. Even in the autumn, when the solar panels generate low amounts of electricity and the heating demands are moderate, they don't need any energy from elsewhere.'

During the project, a simulation tool was developed that enabled the analysis of the scenario above. 'The tool is very useful to explore "What if" scenarios. What happens when you add an extra battery, or when you replace all gas-based boilers with heath pumps? Will the electricity networks of today be able to cope with increasing numbers of electric vehicles in a district?'

Cheering for the pizza test

The researchers conducted various field tests to validate the simulation tool. 'For example, we conducted a stress test together with Alliander in three streets in Lochem: in this so-called "pizza test", we asked about twenty people to simultaneously charge electric cars, and put a pizza in their ovens. Around eight o'clock that evening all of the residents of those streets cheered because their joint actions had resulted in an outage: the fuse in the distribution station had blown, just as our models had predicted,' he laughs. Then, serious again: 'This was a very instructive test, not only for us, but also for Alliander. We discovered, for example, that the three phases of electrical power that are fed into each house, are not equally used. Electricians who can choose freely which phase to connect to were found to prefer one of them over the other two. This leads to non-equal phase loads, which should be measured separately to be able to reliably control the network. Furthermore, it became very clear that there is an urgent need for smart strategies to ensure that the charging of electric vehicles will not lead to power outages.'

Although this project has finished, the work continues, says Smit. 'We are still cooperating with Alliander in a couple of other projects, for example to test different demand-steering mechanisms like variable pricing. They are going to implement some of our ideas and solutions in small towns around Enschede to test their performance in real life. And we are looking into the possibilities to further develop our simulation software and make it available for broader use.'

Researchers

Prof. dr. ir. G.J.M. Smit, emeritus
Prof. dr. J.L. Hurink
Prof. dr. C. de Persis
Prof. dr. A.J. van der Schaft
P. Monshizadeh Naini MSc
Dr. T. van der Klauw
Dr. G. Hoogsteen

Users

Alliander NV
Ipsum Energy
Westnetz

Awareness about good and bad practices

Peter van der Sluijs
corporate strategist at Alliander NV

'When we started this project four years ago, the question of robustness and resilience with regard to local demand management was not a topic of conversation amongst energy grid companies. All problems had been solved with grid investments, even if these were not economical. The CPS programme has created a broader awareness amongst governmental bodies and network providers that there is a right way and a wrong way of implementing smart grids and smart markets.

For us it is important to get an idea of how different scenarios will work out in practice. Is a specific problem expected to occur incidentally or will it be of a more structural nature? Should we act now, or can we take time to investigate possible solutions? Can customers be incentivised to prevent or correct potential problems? In that respect, the field test in Lochem was very instructive. It was reassuring to experience that the current grid is pretty robust. We had to fire up a lot of ovens, heaters and electric cars to create a demand spike that caused an actual outage.

Digitisation of the electricity grid introduces new solutions to reduce the cost for customers. It also creates new dangers that we should be aware of and prepare for. As a result of this project, we have gained valuable insights about how to analyse, predict and hopefully avoid problems. We have already incorporated some of the findings in our grid management practices. For other insights to be useable, we need to amend existing laws and regulations. This will take time, but the results of this project certainly help to convince lawmakers about the benefits of amendments.'

Mastering simulation for virtual prototyping

Prof. dr. J.J.M. Hooman, Ir. C. Gortemaker **project** 12701

Almost every part or function of a complex high-tech system can be simulated. But all of these simulations run separately, and often provide information about their own specific function only. In this project, researchers from Radboud University and the University of Twente worked on a prototype of 'the master of all simulations': a model-based methodology which combines all available simulations into one replica of the entire machine.

As soon as project leader Jozef Hooman starts to talk about the ultimate goal of this research project, one cannot help but wonder why his ultimate solution does not exist yet: 'We are looking for ways to connect different simulation tools to each other, in order to make a complete virtual prototype of a hardware system that is controlled by software.'

Traditional practice in high-tech production is often as follows: first the mechanical parts are fabricated, then the electronics are implemented, and finally the software is written to control it all. 'And then the most stupid errors can occur while testing this prototype,' says Hooman. Sensors that cannot measure the desired parameter due to some unexpected interference. Software control loops that force the entire machine to freeze into a certain state when a single sensor breaks. Or even disastrous collisions of moving parts because of some minor fault in the electronics.

Smarter prototyping

'Manufacturers would like to know at an early stage how choices in the design phase influence the performance of the actual system, without having to build a costly and time-consuming prototype first,' explains Hooman. The use case for this project was provided by Malvern Panalytical, and involved their X-ray diffractometers. 'They would like to also produce modified versions of this system to serve new markets with. But before they start producing these, they want to get a good idea of how changing the dimensions of the system will influence its resolution and overall performance.'

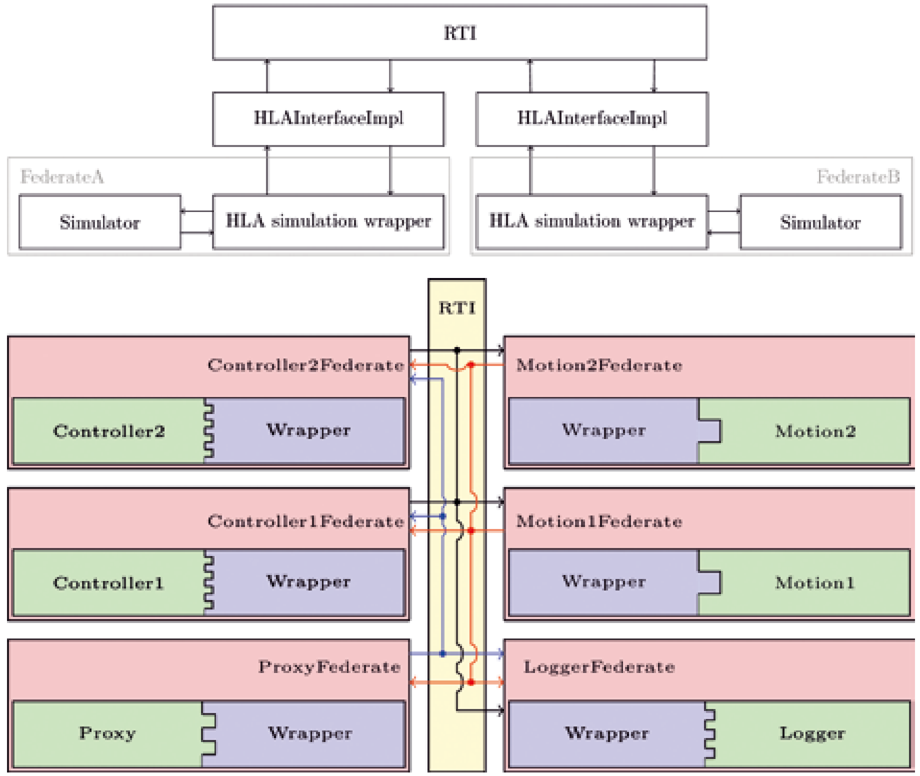

The software structure as developed within the project.

The desired virtual prototype can only be constructed if all available models describing the behaviour of the different parts of the system are combined into one single simulation tool. 'For the hardware, you can start with the CAD drawings. First, we determined to what level of detail we needed these drawings to be. We don't need information about every bolt or joint, as long as we can mimic the overall performance of the machine accurately and fast enough.'

Master organiser

Since simulation tools come in all sorts and sizes, leading to a wild variety of information types and data formats, an automated overarching system is needed to tie them all together. 'We use the standard high-level architecture (HLA) for these purposes,' explains

Hooman. Moreover, to interact with simulation, the Functional Mock-up Interface (FMI) standard is used. This open format for exporting and importing simulation models allows for a high level of automated processing. ‘We only need to feed the master layer with information about the connections between the different simulations, such as what types of information each individual simulation needs and provides, and which processes influence each other in what way. Then the rest of the software connecting all the inputs and outputs is generated automatically. The HLA acts as a supervisor, organising time slots and sequences based on the conditions set by the user.’

Since this methodology is generic, it can be used for all kinds of co-simulation challenges. ‘During the course of the project, we have not only tested it with the diffractometers at Malvern Panalytical, but have also applied it to a case study of a smart lighting system in the Witte Dame in Eindhoven. There we simulated automatic lighting of 90 rooms and their connecting corridors. And to see if our methodology was scalable, we successfully ran it in the cloud on multiple nodes in parallel.’

Besides the obvious advantage that you can simulate the performance of an entire system before it is built, virtual prototyping has another major advantage too, says Hooman. ‘You can use it to mimic what happens in the case of an error. Imagine one of the movement sensors in an expensive machine like a medical X-ray scanner malfunctions. I’d rather see the scanner slam into a virtual patient on a screen than witness that in real…’

Researchers

Prof. dr. J.J.M. Hooman
Dr. ir. J.F. Broenink
T.C. Nägele MSc
T.G. Broenink MSc

Users

Controllab Products BV
Malvern Panalytical
Océ Technologies BV
TNO-ESI

Surprised about feasibility

Chris Gortemaker
group leader hardware Malvern Panalytical

‘We would like to get a complete virtual model of our systems to be able to see how a transformation to different scales and materials will affect their performance. In this project, the first steps towards such a virtual model have been made. The software layer is more or less finished and ready for implementation. We are now waiting for the hardware part to come through.

Perhaps things might have gone faster if besides the two PhDs and the PDEng some master students had been involved who could have explored the more practical implementation of the devised models. For example, one thing we desperately need is a practical way of working to automatically translate CAD drawings of the systems into physical parameters like vibrations, moments of inertia and so on. This would also be very useful for others, since it is a generic problem.

This project has triggered us to further invest in developing virtual models that can prevent underperforming prototypes from being built. I could have well imagined during the project that the whole idea of a virtual system would have proven to be unfeasible. But to my surprise, after three years of serious efforts, it still seems to be a feasible goal.’

Colophon

Final Report Perspectief Programme Robust design of Cyber-Physical Systems

Publisher
Netherlands Organisation
for Scientific Research
Domain Applied and
Engineering Sciences
P.O. Box 3021
3503 GA Utrecht
The Netherlands
+31 (0)30 600 12 11
ttw@nwo.nl
www.nwo.nl

July 2018

Text and interviews
Sonja Knols, Ingenieuse

Co-ordination and production
NWO-domain Applied
and Engineering Sciences

Photography
Institutions involved
Shutterstock: page 4, 23, 39,
46-47 and inside cover

Cover photo
A platoon of test vehicles. Photo: TNO

Design and realisation
Argante Argante, Amsterdam

Print
Zwaan printmedia, Wormerveer

FSC
www.fsc.org
MIX
Papier van verantwoorde herkomst
FSC® C014218

Publisher
Netherlands Organisation
for Scientific Research

Address
Van Vollenhovenlaan 661
3527 JP Utrecht

Mailing adress
P.O. box 3021,
3502 GA Utrecht
+31 (0)30 6001 211
ttw@nwo.nl
www.nwo.nl

NWO | Applied and Engineering Sciences