Learning in control

Citation for published version (APA):
Oomen, T. A. E. (2024). *Learning in control*. Technische Universiteit Eindhoven.

Document status and date:
Published: 14/03/2024

Document Version:
Publisher’s PDF, also known as Version of Record (includes final page, issue and volume numbers)

Please check the document version of this publication:

• A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
• The final author version and the galley proof are versions of the publication after peer review.
• The final published version features the final layout of the paper including the volume, issue and page numbers.

Link to publication

General rights
Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the “Taverne” license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy
If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

Download date: 17. May. 2025

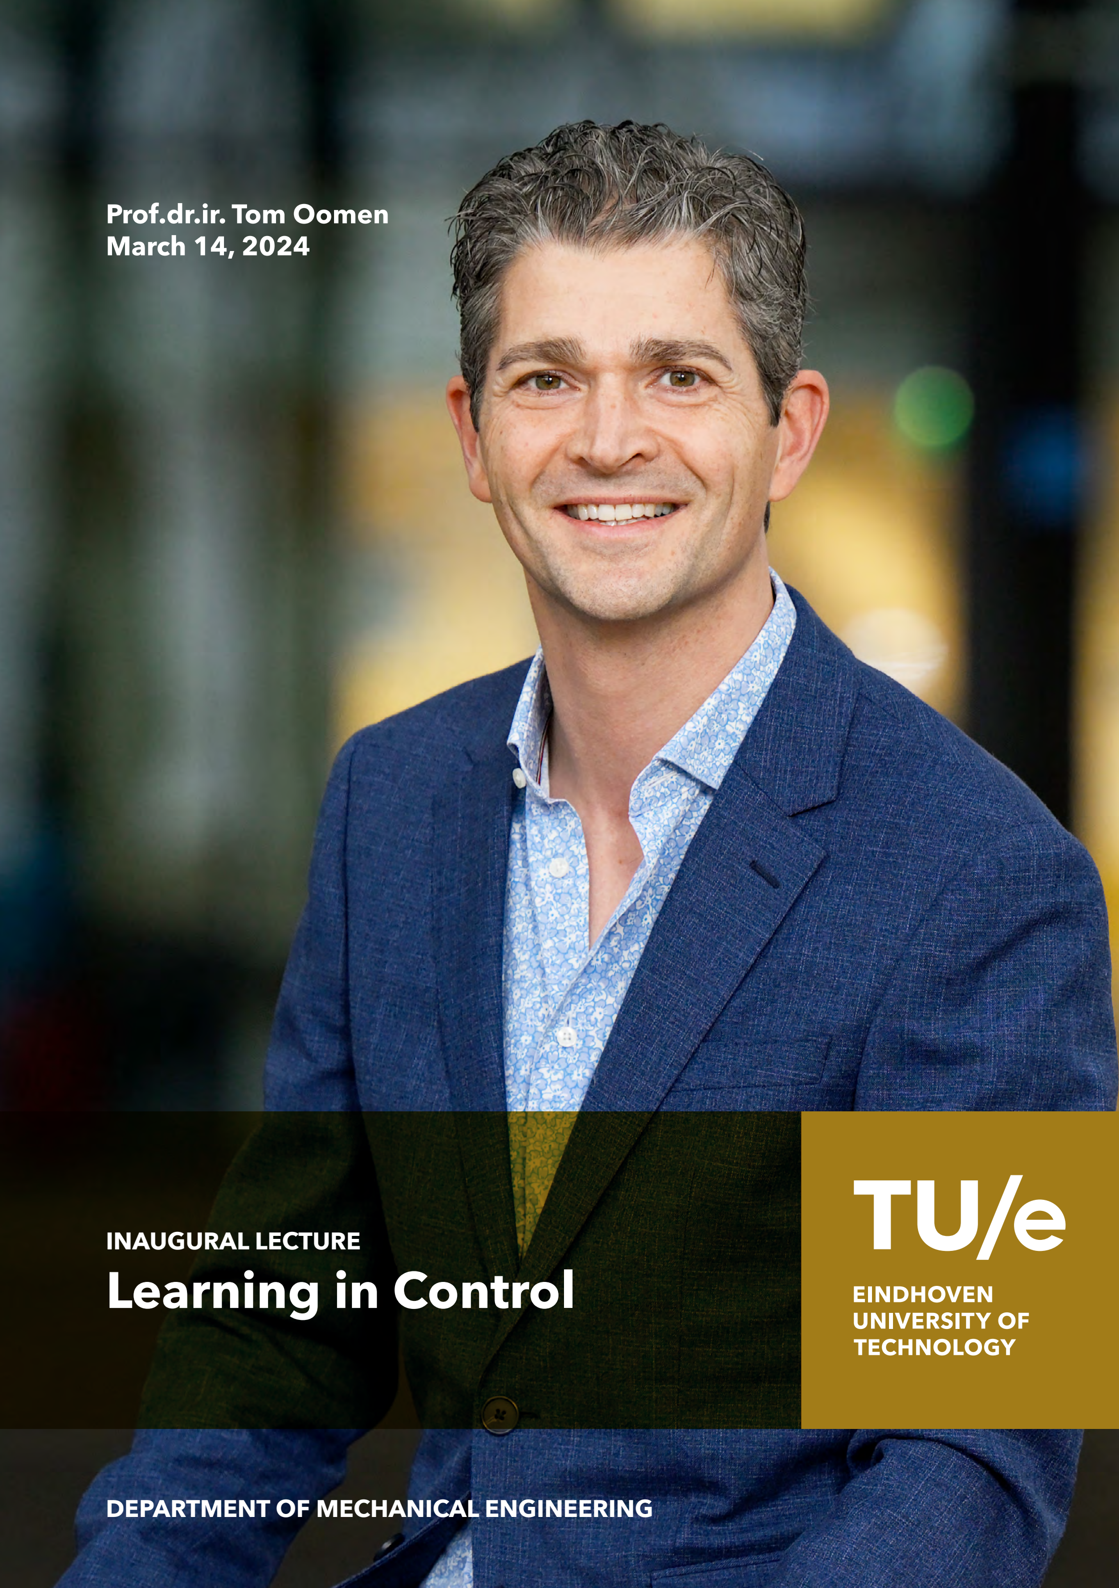

Prof.dr.ir. Tom Oomen
March 14, 2024

INAUGURAL LECTURE

Learning in Control

EINDHOVEN
UNIVERSITY OF
TECHNOLOGY

DEPARTMENT OF MECHANICAL ENGINEERING

INAUGURAL LECTURE PROF.DR.IR. TOM OOMEN

Learning in Control

March 14, 2024
Eindhoven University of Technology

In control

Dear Rector Magnificus, colleagues, friends, and family,

Almost three years after my formal appointment as a professor, it is a great pleasure to welcome all of you to my inaugural lecture. The most important aspect in preparing any lecture is to consider the audience. When sending out the invitations, I realized that the audience is very broad, with three groups: academia, industry (the largest part), and a large group with a general interest and diverse backgrounds.

One of the main questions that many of you have asked me during all these years is “What do you do?” Typically, this is a conversation that goes as follows. I reply, “I work at Eindhoven University of Technology, doing research, teaching, and so on.” The next question already gets a bit tricky and is typically, “what field?” Then I typically reply with “mechanical engineering.” The typical response then is, “like welding, milling, and so on?” Of course, the response is understandable, but completely wrong, so I reply then, “no, I am actually in control; think of machines that have to position something and you have a sensor that looks at the current position and then an algorithm has to compute a current for the motor.” In many cases, this leads to something along the lines of, “oh like making software?” Me: “no, no…” Today, I will share with you the beautiful hidden technology called control that is intimately connected to mechanical engineering.

Now you may think that such questions are only asked by people who did not study engineering, but this is not the case. In fact, another interesting question that surprised me and caused me to reflect actually came from an academic colleague. This person, also in the field of control, asked me, “Why did you study Mechanical Engineering? I had academic ambitions, so I studied…” Interestingly, I actually ended up studying Mechanical Engineering by coincidence. About 25 years ago, I decided to pursue a study that relates to technology, so I attended open days at multiple universities and broad study programs, including Electrical Engineering, Aerospace Engineering, Applied Physics, Computer Science, and Artificial Intelligence, and I joined another student who was interested in Mechanical Engineering at Eindhoven University of Technology. Back then, the unexpected visit to Mechanical Engineering inspired me most. Currently, I feel that the Mechanical

Engineering departments where I work in both Eindhoven and Delft are ideal for the fundamental and academic work I do. As you will see today, I now interact with applications in all the disciplines I visited 25 years ago, which is a lot of fun and very inspiring and fits naturally with what I enjoy.

Today, I will share with you what I have been doing but mostly what I want to do as a professor. But also, what I enjoy doing, so let me start with what has fascinated me for many years.

Technology, science, and knowledge

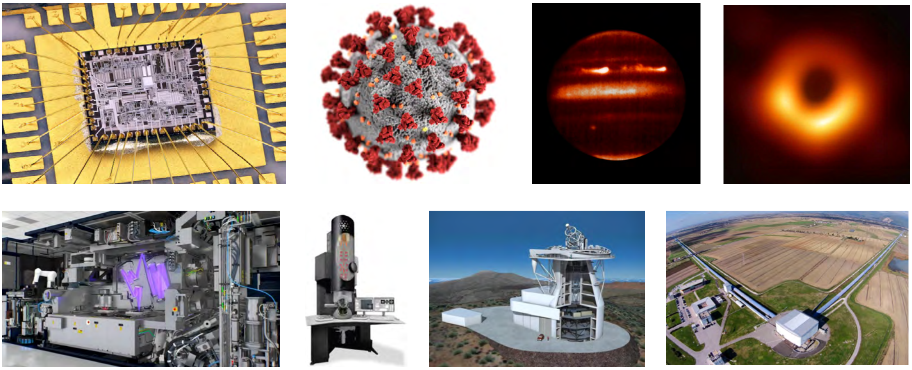

Figure 1. Technological and scientific developments and the enabling mechatronic systems. Top row from left to right: integrated circuit, image of a virus, image of a planet, image of a black hole. Bottom row from left to right: wafer scanner for producing integrated circuits, electron microscope for imaging at a small scale (such as viruses), telescope for exploring the universe, gravitational wave detector that has measured the merging of two black holes.

Looking back, and some of you have known me for a long time, I have always loved technology. I believe we live in one of the most interesting times so far and I am very happy to be born in this time. Let me take you back approximately 30 years, since I still have a few memories that are very sharp. I was about the age of my eldest son, reading in a magazine that in the year 2000 – which seemed very far away back then – everyone would have a mobile phone. I remember I felt that this was complete science fiction. Personal computers were not part of many homes and a lot of people thought they would never need one. Today, a mobile phone is much more powerful and connected than a home computer of the 1990s and most of you will probably check your phone multiple times during this lecture. This huge development is all due to the development of chips, see Figure 1, in which our region has a key role.

A second example of the developments we have made can be found in the recent pandemic. Despite the negative impact, I believe it is extremely remarkable how quickly researchers obtained an image of the virus, see also Figure 1, within just a few weeks after the first symptoms occurred. This is truly amazing, especially if you realize that humanity only discovered that we get sick through viruses relatively recently in our entire history, about a century ago.

Last but not least, many of my generation enjoyed the large number of science-fiction series and movies that were popular. These often took place in space and often featured other planets, as well as very mysterious and intriguing black holes. Besides visualizing planets and galaxies, researchers managed to produce an image of a black hole a few years ago.

A key technology behind all these developments is mechatronics and control. This is the field I work on in my research. Let me briefly take you through the mechatronic systems that relate to the earlier examples in Figure 1.

The first one is a wafer scanner, which is probably the most expensive industrial mechatronic system available. This is the most unique and essential machine in the world that is critical to the production of integrated circuits, or chips. This machine is developed in the Eindhoven region by ASML [1].

In the second example, you see a virus that was imaged by an electron microscope of Thermo Fisher Scientific, another company based in Eindhoven. Such electron microscopes are the key equipment to imaging materials at a very small scale.

The third system is a large telescope, used to produce images of the universe. This ranges from the telescope shown here that is used to image the sun to telescopes for imaging black holes. Telescopes are highly complex systems and I will explain the working principles later in my lecture. Together with TNO, a Dutch institute, we work on the control of this application.

The fourth system shown here also relates to black holes. The setup you see here is a huge interferometer used to detect gravitational waves. These gravitational waves were predicted more than a century ago but had never been directly observed until very recently. Actually, at the time of the invention of this technology, gravitational waves were even assumed to never be measurable from earth. In 2016, the gravitational wave detector you see here was used to detect the merging of two black holes. I will address our work on gravitational wave detectors later.

So far, I have shown you several mechatronic systems we are working on and I will come back to these examples later. I have also mentioned that we work in the field of control, which is a bit hidden in all these mechatronic examples but is essential to their operation. Thus, the question remains: what is control? To explain this, examples that are typically used are the central heating system in a house or cruise control in a car. In my view, the field of control is much more exciting and it has a very rich history. Actually, control takes us back much longer than you may be aware of.

A history of control: mechanical engineering

The origins of the field of control date back a long time and the first applications appear around 270 B.C. Early clocks, especially those used to measure time during the night, consisted of just two buckets with water dripping from one to the other. The major drawback of this design is that if the water level drops, the pressure decreases and the flow towards the bottom bucket reduces [2]. This, of course, complicates determining the actual time, as can be seen in red in Figure 2. In answer to this, Ktesibios devised a mechanism with a float that regulates the height of the water, leading to what we call a linear scale. This enables an accurate time reading, even in the absence of direct sunlight.

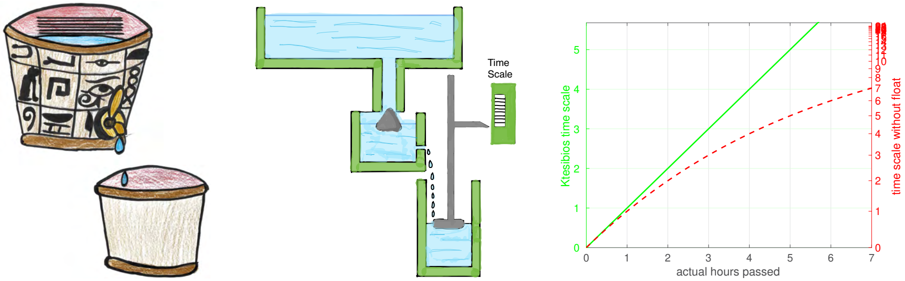

Figure 2. An early example of a control system in a water clock, 270 B.C.

This brings me to the definition of control. For me, an appropriate definition is “control is the art of making systems, typically engineered systems, behave the way you want them to behave.”

Later, control had a major impact in the industrial revolution, which took place in the 18th century. The development of steam engines required control of the engine speed. One ingenious mechanism, called the centifugal governor, has balls connected to the outgoing shaft of the engine. If the speed increases, you can probably imagine that the balls move outwards. Through a mechanism, this pulls a lever that controls the throttle valve. This is actually a refined control mechanism that we nowadays call proportional control. Interestingly, if anyone ever tells you he or she is going balls-out, it actually originates from this control application, meaning “with maximum effort, energy, or speed and without caution or restraint.”

Mechatronics, integration of disciplines

Mechatronics is a relatively new discipline. Not too long after the Second World War, many developments took place. In Japan, machine developments led to a concept called Mo-Chin-Trol, originating from 'Motor+Machine+Control' = Mo-Chin-Trol. Interestingly, control was an explicit part of the name. In that period, the controls were mostly analog electric circuits. The controller then evolved towards electronic implementations for more versatile tasks. This led to the current name of mechatronics, combining 'mechanics' and 'electronics' [3], applied for by Yaskawa Electric Corporation in 1969 and published in 1971. Although hidden from the name, control remains an integral part of the system.

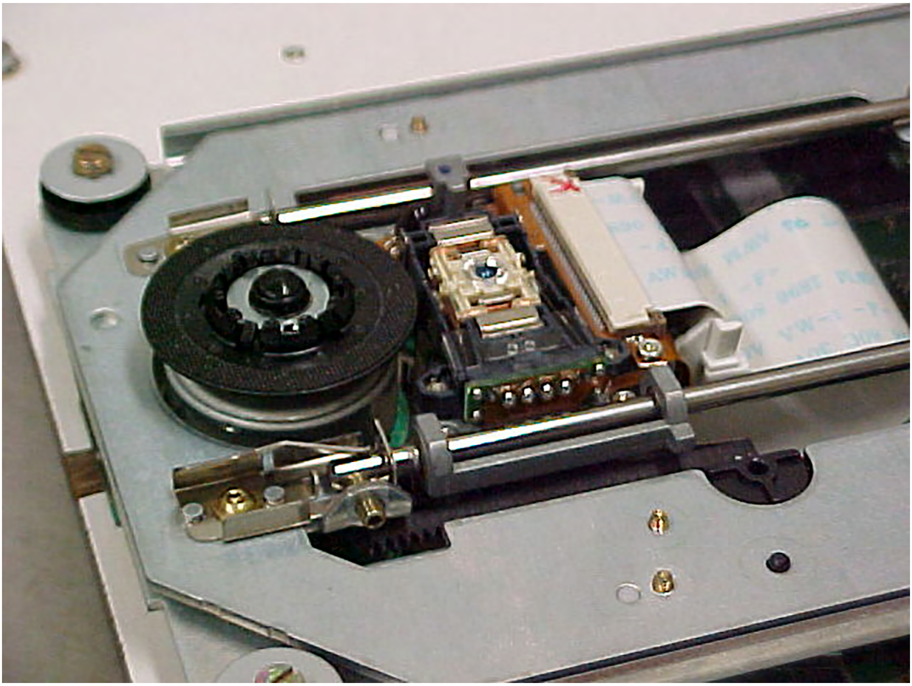

Figure 3. Optical disc drive.

A successful example of a mechatronic system that most of us still use is the optical disc drive, which was co-invented by Philips [1], see Figure 3. Here, a laser points at a disc. The reflection of the laser contains encoded information, which is stored by dots on a track. At the same time, the reflection also contains information on the position of the laser with respect to the track. Based on this position information, a control algorithm determines the appropriate inputs for the electric motors to position the lens. The main goal is to keep the lens in focus and positioned radially with respect to the tracks.

Two fields have had a major impact on control. First, the controller is implemented in electronics and software, which has been a key enabler of more complex algorithms and inexpensive production. Second, the field of mathematics has been substantially developed and a complete theory has been developed to design general controllers with stability guarantees [4].

My first personal encounter with control in a real-life system involved this optical disc drive during my internship at Philips Optical Storage in 2004. I applied control design approaches, which were refinements of the methodology in [4], to a real-life system with all its aspects, integrating it with experimental modal analysis techniques. Despite the fact that I received an excellent grade and the company was satisfied with the results, I also felt the work was mostly applying what I learned during university courses and had a desire to continue on a problem that had a stronger scientific part. I therefore made up my mind to go back to the university for my MSc thesis research. Several interesting things happened in the spring of 2004. First, during the laudatio of a PhD candidate for whom I did a research project, Maarten Steinbuch mentioned that “Geert can do a PhD but does not want to do it. Tom will do a PhD but does not know it yet.” The idea of pursuing a PhD had never crossed my mind. Second, Maarten sent me to Okko Bosgra, his former PhD advisor, who now had a part-time professor position in the group. Okko actually sent me back to Philips, but now to a different building at the same industrial estate, Strijp-S, where the Philips Centre for Industrial Technology was located. The assignment involved precision motion control, working on arguably the most advanced industrial mechatronic system on the market: a wafer scanner.

Wafer scanners are the key technology behind Moore’s law, which dictates a doubling of transistors on chips roughly every two years. This prediction was made in 1965 by Gordon Moore, who was a co-founder of Intel [5]. This exponential growth gives chips the enormous capacity they have today. To make exponential growth tangible, let me go back to the earlier example. 30 years ago, our home computers had approximately 275,000 transistors. Currently, mobile phones have processors that contain 15 billion transistors, which is an increase of more than 50,000 times. Interestingly, your mobile phone is also much smaller than the old home computers. The way in which this exponential growth of transistors is achieved is through shrinkage, meaning that the components in chips are made smaller and smaller.

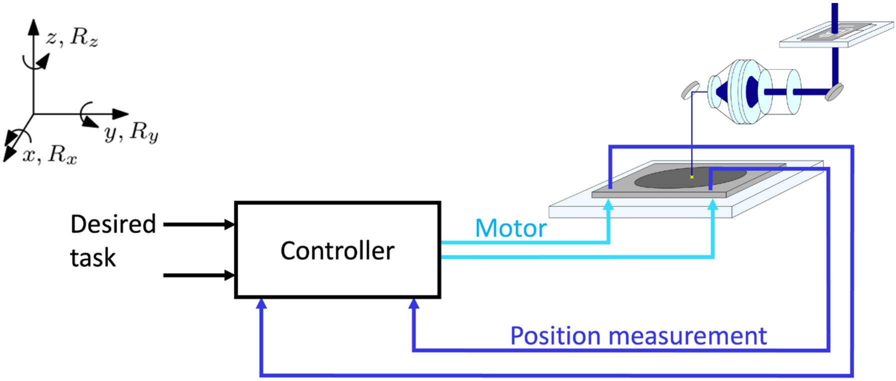

Figure 4. Motion control in six motion degrees of freedom.

One of the key subsystems of a wafer scanner is a wafer stage, which contains the wafer. On this wafer, chips are eventually produced through a photolithographic process while the wafer moves with a constant velocity. The massive growth of computation power has two conflicting requirements of wafer stages. First, the shrinkage of components means that the wafer must be much more accurately positioned, in the order of nanometers. Second, throughput and thus market position are improved by increasing velocities and accelerations. These requirements imply that the fast motion of the wafer stage must be controlled to extreme accuracies in all six motion degrees of freedom: three translations and three rotations.

The motion control system of the wafer stage employs position measurements with sub-nanometer accuracy. The controller compares the position with the desired trajectory and ensures that the actual position does not deviate more than a few nanometers from the desired trajectory. This is achieved by continuously selecting an appropriate input current to the motors, as illustrated in Figure 4.

Currently, state-of-the-art mechatronic systems, including wafer stages, are designed through a relatively stiff mechanical design. This avoids any mechanical deformations; as a result, the six motion degrees of freedom can be controlled independently, essentially through traditional control designs that build upon the methodology in [4].

Having introduced the system, let me now introduce my research assignment. A major challenge that arises due to the extreme accuracy is at the heart of mechatronics, i.e., the integration of mechanical motion control and electronic implementations. Indeed, as already anticipated in [4], electronic implementations would require a process called sampling, involving only measurements at prescribed times. The underlying position of the mechanics evolves in continuous time. Combining these in a way that has mathematically provable guarantees, while at the same time being compatible with control engineering methodologies, involved a major theoretical challenge and led to my MSc thesis [6]. This also sparked my application-driven fundamental research line, involving fundamental research from problems occurring and with direct impact in real-life applications [7].

Let me zoom out again and summarize the role of the field of motion control. Motion control is essentially the brain of the machine, which enables fast and precise motion. Indeed, a striking similarity with nature occurs: mobile organisms, such as many animals, have brains. Organisms that have zero or limited mobility, such as plants, coral, and sponges, do not have brains.

Envisaged future data-intensive machines

My MSc research triggered my interest in application-driven fundamental research and, in the years that followed, I continued pursuing a long-term dream that is still very relevant. Let me start by looking into the future and later highlighting a few recent developments.

The increasing requirements on accuracy, speed, and versatility, driven by Moore's law, impose requirements on mechatronic systems. The design of such mechatronic systems is nowadays truly multidisciplinary, involving physics, optics, heat and flow, materials, mathematics, and software [8], [9]. In 2000, a renewed definition of mechatronics redefined mechatronics as [10] "the synergetic integration of physical systems with information technology and complex-decision making in the design, manufacture and operation of industrial products and processes."

To achieve breakthrough performance to meet the increasing demands from society, I envisage a radically new lightweight system design to achieve unparalleled speed. Indeed, according to Newton's second law $F = ma$, reducing the mass of the system directly leads to an increased acceleration. Reducing the mass comes at a cost, since this typically leads to a situation where the stiffness reduces and the system has pronounced flexible mechanical behavior.

The concurrent development of ubiquitous computing, together with new sensors, actuators, and connectivity, provides a major opportunity for data-intensive future machines in which active control is used to compensate flexible dynamical behavior towards ultimate overall system performance. This will lead to future self-adapting, improving and self-healing high-tech systems.

Currently, we are taking several steps towards our future vision, which I would like to outline to make these future ideas more tangible.

First, the idea of lightweight systems with predominant flexible behavior directly leads to a chicken-and-egg problem: without evidence that active control can address the issues arising due to flexible lightweight dynamics, the design of such a system has too many risks due to possible problematic behavior that cannot be controlled properly. To mitigate the risk, we have investigated the control issues in multiple prototype motion systems. In the flexible beam in Figure 5, we have used actuators and sensors at the edges of the beam. These can be used to control the translation of the beam, showing a clear similarity with Figure 4. Continuing the comparison with the wafer stage, the performance is required in the middle, indicated by the variable *z*, where exposure takes place. Under a rigid-body assumption without any flexible deformations, the blue response in Figure 5 is obtained, which is a typical response of a motion system. A temporary measurement, which is only available in the prototype setup, shows a very different behavior: the red response reveals severe deformations. Such a system would not meet performance requirements, e.g., those in lithography [11]. With advanced control algorithms that we are developing, the sensors can be used in conjunction with a model of the mechanics to estimate the spatio-temporal deformations of the beam. This can be used to use the existing actuators to control the performance variable, as I will later show.

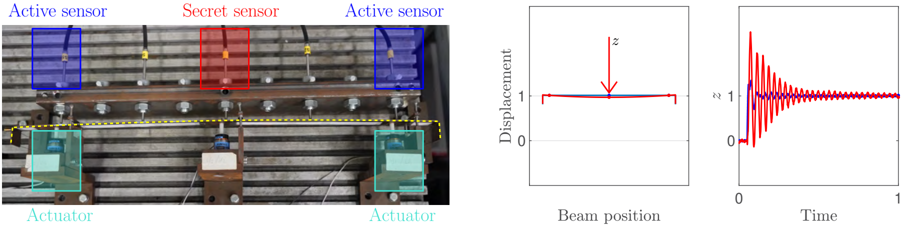

Figure 5. Flexible motion system prototype and expected performance with traditional design methods (blue) and achieved performance due to flexible dynamics (red).

Second, multiple actuators can be used to increase stiffness and damping through active control. The availability of actuators that can generate a detailed spatial force profile is a standard feature of new motor concepts but is not actively exploited for this purpose. We have applied this idea of over-actuation to a prototype wafer stage to actively control the torsion mode. This leads to a much larger bandwidth and performance of the system [12]. A key challenge in this approach is the complexity associated with the large number of actuators and sensors that exhibit strong interaction. Hence, traditional control design approaches cannot be applied.

Third, traditional mechatronic designs divide the overall goal into manageable and monodisciplinary subproblems that introduce unnecessary performance limitations. We envisage a radically new centralized control approach, where all subsystems work together to achieve the overall system performance, see Figure 6. This includes [13] a synchronization of wafer and reticle stage, the active control of thermo-mechanical deformations, and adaptive optics.

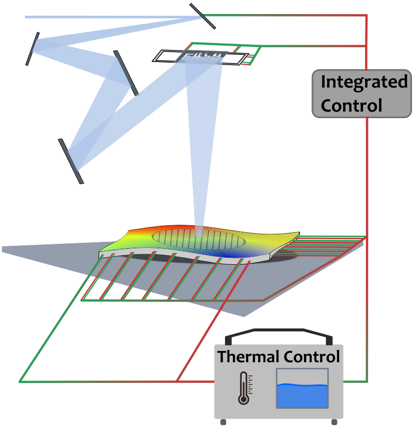

Figure 6. Overall machine control: the pinnacle of data-intensive future machines.

Evaluating these envisaged developments reveals that the last example is the pinnacle of what we envisage as a data-intensive future machine, where a massive number of inputs and outputs is used to achieve optimal overall machine performance. The major roadblock is the complexity: such a mechatronic system involves a complexity that is far beyond what state-of-the-art motion control approaches and algorithms can deal with. The main opportunity is the large availability of data that can be exploited. Our goal is the development of a general control methodology to control real machines to the limits of performance by learning from data.

Before going deeper into the opportunities, let me emphasize that the specific application is one example of a complex motion system and one I have worked on for a long time. Our goal lies in the development of a general methodology for a range of different applications. I will first outline the complexity associated with several applications we work on.

First, large flatbed printers are used for industrial printing. The main difference compared to the printer you may have at home is that the media does not move, but the carriage that contains the printheads moves over the media in two translational and one rotational motion degrees of freedom. The moving mass of the carriage leads to position-dependent dynamics, leading to highly complex dynamics and control algorithms.

The earth-based telescopes in Figure 1 are used to image the sun, exoplanets, and distant galaxies. A key complication is that the wavefront of the received light is heavily distorted by the earth's atmosphere. A prime example of this distortion is the twinkling of stars: what you are seeing is due to atmospheric distortions.
By measuring the wavefront and using a deformable mirror in the optical path, the wavefront can be actively compensated to obtain sharp images of the universe. Next-generation deformable mirrors have thousands of actuators and lead to a massive increase in the complexity of the control problem.

A last example of a highly complex opto-mechatronic system is the gravitational wave detector in Figure 1. This system is essentially a large interferometer with arms of 3 km in length. Due to an extremely precise design, it can measure a relative change in arms of 10^{-18} m due to a gravitational wave, e.g., when two black holes merge. The main complexity lies in the extremely accurate control of multiple subsystems with dominant disturbance sources to achieve the overall system accuracy.

These applications share a large complexity and extreme performance requirements. Next, I will outline how new control methodologies that exploit the ubiquitous data are foreseen to be a key enabler for these systems.

Learning models for control

Control design requires a model of the system that describes how the system will behave when we apply a certain input. There are essentially two ways to model a system. First, one can model the system using physical laws. This is typically a rather time-consuming process. While physical modeling can provide a deep understanding of how the system works, it is often too inaccurate to use for control purposes. The second approach is to learn models from data, a process also called system identification.

System identification involves a few basic ingredients. First, an experiment is designed and an input is applied to the system's actuator. Then, the output is measured using the sensors. The second ingredient is model parameterization. The model parameters are then optimized such that the model output resembles the measured output in an appropriate criterion.

Among system identification methodologies, frequency domain approaches are the bread and butter of motion control engineers. These frequency response functions are used directly as a basis to tune controllers using a process called loop-shaping. The process of obtaining these frequency response functions has been developed substantially since the early developments more than half a century ago. In the 1990s, the use of multi-sine input signals was strongly advocated to improve accuracy. However, these lead to a prohibitive measurement time for future mechatronic systems with a massive number of inputs and outputs.

A major step in identifying frequency domain models for motion control is the incorporation of prior knowledge. Basically, we know that the system involves mechanics with characteristic behavior in the frequency domain. We exploit this by estimating local parametric models. By developing such methods, we have been able to reduce the experimentation time of a system with eight actuators and six sensors from one afternoon to 1.5 seconds. Recently, we have extended these results to a deformable mirror for a telescope. This deformable mirror has 52 inputs and 26 outputs. Due to the extreme complexity, only static models are traditionally estimated. However, these static models are insufficient for high-performance control in dynamic ranges. Using our new methodologies, we can now very accurately identify all 1352 transfer function elements using only one minute of measurement time.

We have very recently taken the fact that we know a lot about our systems a step further. In particular, the deformations behave in a specific manner since these take place in the mechanical domain. In the earlier example of the deformable mirror, this allows us to extend the model to include an additional 26 outputs without having measured at these specific locations on the structure. This paves the way to design inferential controllers that predict and control the performance on the point of interest where performance is required but cannot be reached by sensors. Initial experimental results on the prototype motion system of Figure 5 are very promising, recovering performance on the performance location z without measuring that signal in the feedback loop.

In summary, the development of new methodologies to learn dynamical models from data through system identification for mechatronic systems is highly promising. Further developments relevant to progress the application domain include the modeling of interconnected multi-physics systems, the development of approaches that focus on control-relevant models with high numerical reliability, and identification algorithms to complement the models with black-box extensions representing parasitic nonlinearities and addressing sampling aspects.

Learning in machines

An alternative to the more traditional model-based control design approach is to further use data through learning. In particular, the recently renewed interest in machine learning and artificial intelligence due to the ubiquitous availability of data and computation power has led to major achievements in computers that play board games, computer games, and, very recently, language models. The key question that arises is what learning has to offer for motion control in mechatronic systems.

The main difference when comparing these simulation environments with mechatronic systems is that learning must be done fast, since experiments are done in real time and exploration steps to improve performance are typically costly. In addition, any adaptations need to be done carefully, since arbitrary actuator inputs may easily damage the system.

The main promise of the field of control is that whatever behavior can be predicted can be compensated for. Typically, a large part of the errors is repeating, for instance due to traditional performance limitations of feedback systems. As an example, repeating ten experiments of the same task often reveals a very similar error profile. This can be completely compensated for.

Our solution lies in integrating data and model-based approaches. In Figure 7, the result of such an approach, called iterative learning control, is depicted. These algorithms ensure fast and safe convergence by using models and their uncertainty, leading to perfect performance up to the sensor resolution in a few iterative experiments under normal operating conditions.

The main limitation of these learning algorithms is that the resulting models are signal-based with limited interpretability and generalization to other tasks. In our current research, interpretable models are being learned, including parametric low-order physical models, as well as flexible models that are extensively used in machine learning such as Gaussian processes and neural networks. We have recently shown that these models are extremely powerful in semicon backend equipment, where these have led to improved motor calibration methods and position-dependent snap feedforward parameters.

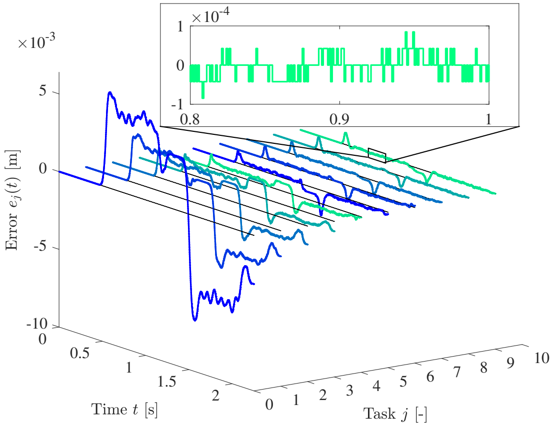

Figure 7. Learning in machines.

The impact of these methods reaches far beyond precision mechatronics; two applications in the medical and energy domain include the following. In mechanical ventilation of patients in intensive care units, breathing also follows a highly repeating pattern. However, the system uncertainty is much larger due to variability in patients, ranging from infants to adults. Our robust algorithms substantially improve the breathing support systems while dealing with safety due to the integration of model knowledge and data. Furthermore, wind energy systems, ranging from turbine control to wind farm control, require overall system performance optimization in a situation with large uncertainties and operational constraints. Here, such data-driven control algorithms are an essential step towards future energy systems.

Learning faults: self-healing machines

The economic value of all these applications is largely determined by their productivity, which is directly related to the system uptime. Without maintenance, it is not a question of whether a machine will fail but rather when it will fail. Traditionally, two approaches can be distinguished as maintenance policy, see Figure 8. In reactive maintenance, the system breaks down, requiring a service engineer to repair the system, typically leading to a substantial downtime. In contrast, in preventive maintenance, systems are repaired well before they break down. Preventive maintenance typically requires less repair time but is performed more often.

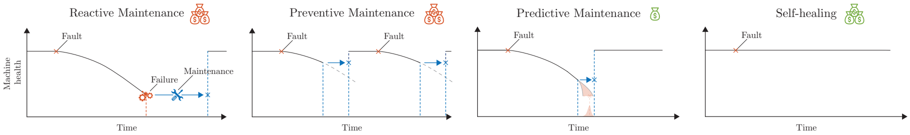

Figure 8. From reactive and preventive maintenance towards preventive maintenance and self-healing machines.

A major opportunity lies in re-purposing the models for fault diagnosis. The predictive power of these models can be employed in a digital counterpart that uses real-time actuator and sensor data. Following the recent developments in advanced motion control, we are currently developing new maintenance approaches. By employing accurate dynamical models, highly accurate fault diagnosis systems accurately predict when maintenance needs to be performed in an optimal way, which is referred to as predictive maintenance.

A major step towards self-healing machines is to integrate the fault diagnosis system with advanced motion control concepts such as over-actuation, which I introduced earlier. If certain components fail, control can reconfigure the actuators to enable self-healing machines, reducing maintenance towards a minimum.

What can we still learn?

Let me shift focus at this point to the role of academia in a rapidly changing world. In my view, the research I have shown fits into application-driven fundamental research, leading to general methodologies with solid foundations on the one hand and solutions to practical applications on the other. Such technological developments are taking place increasingly fast in an increasingly connected world. Indeed, many developments are a result of the earlier mentioned Moore's law, which is exponential. What are the implications for the positioning of academia, the knowledge machine, and the use of that knowledge? There are many aspects related to this, so let me address a few concrete points.

PUBLISHING

The first aspect I want to touch upon is how knowledge is being disseminated in an academic environment: the publishing of papers. There are many aspects related to publications, so let me start where it typically starts in practice: with conference publications. Conferences are the places where results are first published and discussed with peers. The most remarkable aspect is that this process has not changed much in the field of control over the last decades, most notably in the time it takes. Several decades ago, our flagship conferences had papers with, at most, six paper pages, since these were printed in very big and heavy proceedings, resembling the size of a phonebook. It typically took about a year to get these reviewed; in the past, contributions were even sent out through paper mail and then it was later printed and shipped. One year later, the researcher would present the idea to the broad community. The most remarkable observations are that 1) it still takes about a year to get this whole process from idea to presentation, despite the fact that reviews are now obtained through electronic systems; 2) publications are now digital PDF files but are still limited to six pages, with an unchanged overlength fee of USD 200 for an extra page. This means that it is still a very slow process and, when researchers attend a conference, they are discussing ideas that are at least one year old.

There are also aspects that have changed a lot. Using today's technology, preprint servers have appeared in the academic world. A typical example is arXiv, with which one can timestamp a manuscript and distribute it even before the conference takes action to send it to reviewers. It is now considered acceptable to cite these preprints, so some researchers are accelerating, but our conferences where we meet and have discussions involve results and ideas from at least a year ago.

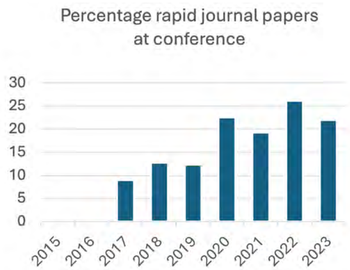

2023	Days until first decision	Days until published
Conference CDC	108	257
Rapid-publication journal L-CSS	31	92

Figure 9. Left: percentage of rapid-publication journal papers in IEEE Control Systems Letters (L-CSS) as part of the flagship conference IEEE Conference on Decision and Control. Right: 2023 publication times for a conference vs. rapid-publication journal.

Luckily, there is also progress. In 2017, I joined the editorial board of a new journal that publishes letters, essentially with the format we are used to in conferences. The idea is rapid publication, whereby we achieve a first decision in 31 days on average. Compared to conference publications, this is a reduction by a factor of three! These journal papers can then also be presented at our flagship conference, CDC, and we have grown to provide 20-25% of the conference material, see Figure 9.

A natural question we could ask ourselves is why we still keep the old-fashioned system of slow peer-reviewed conference publications for the other 75% of the conference. We could use our conference to discuss recent results, abandon the overlength paper fees now that everything is digital only, and also abandon reviews of such material to reduce our community review duties by a factor of four.

Let me mention two more aspects on this side. First, the number of papers appearing is increasing. This is in part due to the increase in academic institutions and their size globally, but also largely due to the increased productivity of authors; in fact, some authors publish a new paper every five days [14]. I believe that a huge amount of information is inevitable in the information age we live in. So, I agree with the point in [14]: “The number of papers should not really count as positive or negative.” The main question is where we want to focus our review efforts as a community, especially since three to five individuals are needed for each paper.

Second, there are multiple workshops in our field that do not have any papers associated with them. Interestingly, the quality of the presentations and interactive sessions is extremely high. Two reasons could be: 1) the focus now lies on the presentation instead of the desire of a researcher to have a paper published; 2) it is recent work instead of ideas from a year ago.

GLOBAL AND LOCAL PERSPECTIVES ON THE ROLE OF ACADEMIA

Let me know shift focus to the practical use of new knowledge. I would like to start with a global perspective. Conferences are a great way to have discussions with peers in academia. At the same time, there has been a long-standing discussion regarding the theory-practice gap in the field of control [15] and there have been plenty of efforts by journals and conferences to appeal to a more practical audience, as reported last year in [16]. A relevant conference – the flagship conference on control applications – is the IEEE Conference on Control Technology and Applications. Last year, in 2023, I served as registration chair, responsible for handling all registrations to the conference. One aspect surprised me: less than 5% of the conference attendees had a non-academic primary affiliation, see Figure 10. This is a small group of people, which mostly consisted of my industrial PhD candidate, invited plenary speakers from industry, researchers associated with government research laboratories, etc. At least from a global perspective, it appears that there may be a substantial theory-practice gap and the impact of our field is rather limited in industry.

Let me shift to a more local perspective on the role of academia in a world that is faster and increasingly global. My main observation is that the theory/practice gap is nonexistent, at least if we look at the local ecosystem around Eindhoven. Before going deeper into this, let me briefly take you back just a few decades. If we look back not too long ago, one of the major tasks of a university was to store books and articles, both in libraries and small offices packed with paper. In particular, many of us, including myself, used to photocopy articles. This has been fully replaced by online services and online search engines. Also, industry has

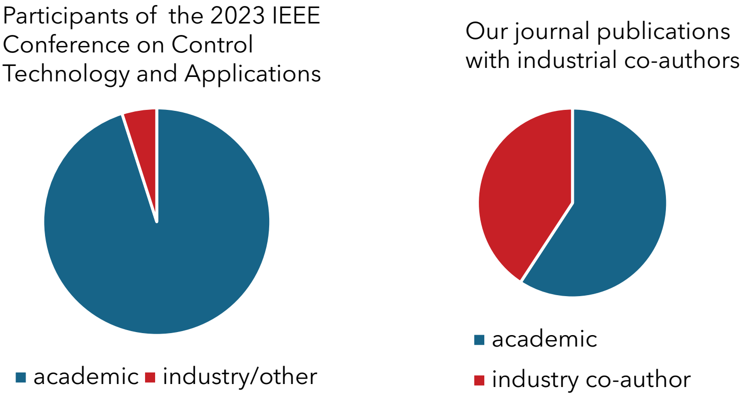

Participants of the 2023 IEEE Conference on Control Technology and Applications

Our journal publications with industrial co-authors

Figure 10. Left: participant distribution of the flagship conference on control applications in 2023. Right: distribution of our journal publications with industry co-authorship.

direct access to all information, even for free with open access publishing policies. Similarly, universities used to have labs to do experiments. In mechatronics, experiments can be done remotely and the physical location of the lab is less relevant. Interestingly, we even had a recent experiment imposed on us during the pandemic and we kept business running as usual to a reasonable extent without even using our buildings. So that brings me to the point: how do we see our role as a university and how do we stay relevant?

Eindhoven University of Technology is ideally positioned in the Brainport ecosystem and has a major opportunity to contribute to major societal and industrial challenges. Many of our researchers do not work in isolation but form a bridge between the global academic community and the local ecosystem. In fact, with several of our industrial collaborators, some of the new research results are implemented much faster than we ever could in academia and, in a recent collaboration with a semiconductor backend company, the ideas of a PhD researcher were immediately adopted and implemented faster than he could do himself. I see this as the best compliment a PhD researcher can have from industry on the relevance of their research and it provides academic researchers with a huge lab with implementation and software support. From a local perspective, the theory/practice gap in the field of control at Eindhoven University of Technology is thus nonexistent. This is confirmed by our journal publications with industrial co-authors in Figure 10, where a very healthy distribution is shown.

A major role of academia is connecting people and several aspects can further strengthen the role of Eindhoven University of Technology in the local ecosystem.

1. Part-time industrial positions. Currently, a part-time assistant professor and an associate professor with a main industry affiliation are in my group. As another relevant example, I recently hosted an industry researcher for one day per week, which led to a very fundamental PhD thesis. Also, all PhD research projects I am involved in typically work on methodology and general fundamental aspects in collaboration with one or multiple partners in our ecosystems. This way, early career researchers learn to pose real-life problems and implement solutions to relevant technical problems, often involving MSc students. In my view, this does not restrict fundamental research but instead enriches it.
2. From a dual perspective, we should embrace part-time positions of TU/e staff. This can be academic and in research institutes, industry, startups, industry teaching/consulting, either in the format of a part-time role or longer sabbaticals.
3. To emphasize, this encourages the embracing of diversity. Not every individual should have the same profile or even have industry collaboration, but it is about the overall community that we want to build.

These aspects are also in line with the broader discussion regarding the transitions towards a fourth-generation university [17].

EDUCATING NEXT-GENERATION ENGINEERS

The most important role of a technical university is educating next-generation engineers. The world around us is changing fast and this directly impacts our teaching. Almost 60 years ago, Bellman [18] stated that “the student in school is trained to furnish precise answers to precise questions. This is hardly adequate training for a real world in which the essence of success is that of obtaining reasonable answers to reasonable questions in areas which are vague and imprecise.” In my view, the current availability of online teaching material provides even more precise answers to even more precise questions. In this respect, we should embrace this, extending the traditional and universal textbooks that dominated education. But the most important role of the university lies in the real questions. We have taken important steps towards challenge-based learning, where students work in groups on topics that intersect our research with societal and industrial demands. An example in our group is the Learning Control course, where students address problems on industrial-scale setups provided by Canon Production Printing and Thermo Fisher Scientific. Furthermore, BSc and MSc projects are often performed in the context of research projects and partially within industrial collaborations.

Closing and acknowledgements

To summarize, I would like to conclude with the observation that we have an exciting future ahead and I am confident that when we look back in 30 years from now, beautiful steps will have been taken that we cannot predict now. I hope to have inspired you on:

- the beauty of the interdisciplinary field of mechatronics,
- the opportunities of learning and control, and
- the opportunities for academia and its ecosystem.

The journey so far has been great and I really look forward to the future. Although I am standing alone on the stage today, this truly was a joint journey and the result of collaborations with many of you. This journey was truly research: we typically have some ideas where to go, but the path is often long, unclear, and with many dead ends. I would like to thank everyone who inspired and motivated me and invite you all to celebrate this together today. I would like to mention a few people in particular, with the risk of forgetting.

First, two people have been mentors to me both early on and for a very long time in my career: Okko and Maarten, who I am grateful for teaching me academic research with ultimate quality standards and the bigger picture, respectively, among endless other aspects.

I have had the pleasure to work with and especially to learn from many inspiring colleagues. In particular, I would like to thank my colleagues at both Eindhoven University of Technology and Delft University of Technology for providing a great environment I love working in. Let me thank three more people who have been around for many years and have been an inspiration and help in countless aspects, Nathan for sharing an office, Maurice for sharing many adventures, and Jan-Willem for sharing a house, among the very many other things.

Thanks to all industrial colleagues from the local ecosystem for the decades of collaboration. I feel very fortunate to have started my career working at Philips CFT, which, after my MSc research, decided to fund my PhD research. Of course, a lot has changed, but it is great to see the ecosystem evolve and continuously see old friends and meet new friends. I learned so much. Thanks!

A special word of appreciation goes to the researchers in the group directly around me, especially all the PhD researchers in the group and all the permanent staff in my group that is currently growing. I have been very fortunate to work with such a team of great people, who are also very talented, creative scientists and engineers. I cannot imagine what an academic career would look like without you!

Also, thanks to all the students: it is great to see the huge interest in our field. Working with students has been a major inspiration during all these years and I am confident that the future is in good hands!

Tot slot wil ik mijn vrienden en familie bedanken. Dank voor alle mooie momenten buiten de universiteit. Mijn ouders wil ik bedanken voor de steun en enorme vrijheid die ze me gegeven hebben. Ellen, Julian, Isa en Mare, dank voor alles, ik vind het prachtig om deze dag met jullie te mogen delen.

Ik heb gezegd

Bibliography

[1] P. van Gerven and R. Raaijmakers, Natlab - Kraamkamer van ASML, NXP en de cd, Techwatch, 2016.

[2] O. Bosgra, "Physical modelling for systems and control, Lecture Notes," Delft University of Technology, Delft, 2005.

[3] T. Mori, "Philosophy of mechatronics," *Journal of the Japan Society for Precision Engineering,* vol. 57, no. 12, pp. 2089-2093, 1991.

[4] H. M. James, N. B. Nichols and R. S. Phillips, Theory of servomechanisms, McGraw-Hill, 1947.

[5] G. E. Moore, "Cramming more components on integrated circuits," *Electronics,* vol. 38, no. 8, 1965.

[6] T. Oomen, "Optimal digital control of high-precision electromechanical servo systems: Concepts and applications," Eindhoven University of Technology, 2005.

[7] D. E. Stokes, "Pasteur's Quadrant – Basic Science and Technological Innovation," Brookings Institution Press, Washington, D.C. , 1997.

[8] H. Vermeulen, "Mechatronica 2.0: Balanceren van complexiteit op componentniveau," Eindhoven University of Technology, 2016.

[9] R. Munnig Schmidt, G. Schitter, Rankers, Adrian and J. van Eijk, The design of high performance mechatronics - High-tech functionality by multidisciplinary system integration - 3rd revised edition, Amsterdam: IOS Press, 2020.

[10] M. Tomizuka, "Mechatronics: From the 20th to the 21st century," *Control Engineering Practice,* vol. 10, no. 8, pp. 807-886, 2002.

[11] H. Butler, "De nanometer onder de duim," Eindhoven University of Technology, 2013.

[12] R. van Herpen, T. Oomen, E. Kikken, M. van de Wal, W. Aangenent and M. Steinbuch, "Exploiting additional actuators and sensors for nano-positioning robust motion control," *IFAC Mechatronics,* vol. 24, no. 6, pp. 619-631, 2014.

[13] E. Evers, "Identification and active thermomechanical control in precision mechanics," PhD thesis, Eindhoven University of Technology, Eindhoven, 2021.

[14] G. Conroy, "Number of 'extremely productive' authors concerns scientists," *Nature,* vol. 625, pp. 14-15, 2024.

[15] D. S. Bernstein, "On bridging the theory/practice gap," *IEEE Control Systems,* vol. 19, no. 6, pp. 64-70, 1999.

[16] D. Y. Abramovitch, "Thoughts on furthering the control education of practicing engineers," *IEEE Control Systems,* vol. 41, no. 1, pp. 64-88, 2023.

[17] M. Bogers and M. Steinbuch, "De vierde generatie universiteit: het nieuwe tijdperk van open innovatie en ecosysteemdenken," *Holland management review,* vol. 208, pp. 63-71, 2023.

[18] R. Bellman, "Mathematical aspects of the theory of systems," in *Symposium of the Symposium on System Theory*, Polytechnic Institute of Brooklyn, 1965.

[19] J. Compter, "Electro-dynamic planar motor," *Precision Engineering,* vol. 28, pp. 171-180, 2004.

[20] E. Evers, M. van de Wal and T. Oomen, "Beyond decentralized wafer/reticle stage control design: A double-Youla approach for enhancing synchronized motion," *Control Engineering Practice,* vol. 83, pp. 21-32, 2019.

Curriculum Vitae

Prof.dr.ir. Tom Oomen was appointed full professor of Advanced Motion Control at the Department of Mechanical Engineering at Eindhoven University of Technology (TU/e) on April 1, 2021.

Tom Oomen studied mechanical engineering and received an MSc degree, specializing in control systems technology (2005) and a PhD degree from TU/e (2010), during which he collaborated closely with Philips Applied Technologies on fundamental aspects in system identification occurring in mechatronic applications. He continued his research on system identification at KTH, Stockholm, Sweden, and at the University of Newcastle, Australia. He was then appointed assistant professor (2012) and associate professor (2018) at TU/e. Since 2021, he has held a joint appointment as full professor at TU/e and Delft University of Technology. Tom and his team pursue application-driven fundamental research in identification, learning and control, with a main expertise in mechatronic systems. Their research is enriched by collaborations with application experts in multiple domains, including medical, energy and scientific instruments. He is a recipient of the 7th Grand Nagamori Award, the Corus Young Talent Graduation Award, and a Veni and Vidi grant. He is currently Senior Editor of IEEE Control Systems Letters (L-CSS) and Co-Editor-in-Chief of IFAC Mechatronics.

Colophon

Production
Communication Expertise
Center TU/e

Cover photography
Bart van Overbeeke
Photography, Eindhoven

Design
Grefo Prepress,
Eindhoven

Digital version:
research.tue.nl/en/

Visiting address
Building 1, Auditorium
Groene Loper, Eindhoven
The Netherlands

Navigation
De Zaale, Eindhoven

Postal address
PO Box 513
5600 MB Eindhoven
The Netherlands
Tel. +31 (0)40 247 9111
www.tue.nl/map

TU/e
EINDHOVEN
UNIVERSITY OF
TECHNOLOGY