European Journal of Control (2007)13:71–81
© 2007 EUCA
DOI:10.3166/EJC.13.71–81

European Journal of Control

In Control, Almost from the Beginning Until the Day After Tomorrow

Jan C. Willems*

ESAT-SISTA, K.U. Leuven, Kasteelpark Arenberg 10, B-3001 Leuven, Belgium

1. A Brief Personal History

I entered the field of control as a Ph.D. student in September 1965, over 40 years ago. A couple of years before, Roger Brockett and Michael Athans had joined the Electrical Engineering faculty of MIT. They introduced *Modern Control Theory*, as the state space approach was called then, as a research activity and had set up a graduate curriculum in this area. In my first semester, I took courses on optimal control à la Pontryagin and on the state space theory of systems. This last course was given for the first time at MIT. On the West Coast of the US, similar courses existed already a few years earlier. Later, I taught this course myself many many times. It became my favorite subject.

After my Ph.D., I joined the control group of the Electrical Engineering department of MIT as an assistant professor. I spent 5 years in that capacity, with a one-year leave of absence as a postdoc at the University of Cambridge in the UK. In 1973, I was appointed Professor of Systems and Control in the Mathematics department of the University of Groningen. In 2003, I became Emeritus. Presently, I am Guest Professor with ESAT (the department of Electrical Engineering) of the K.U. Leuven in Belgium.

During my career, I had the privilege to serve as supervisor to a number of very talented Ph.D. students, 'promovendi', many of whom became important contributors to the field. The number of my Ph.D. students or postdocs at any time was never large, typically 2 or 3. With some of them, I discussed research almost on a daily basis. Working together with these young people has been an important part of my *modus operandi* of doing research.

I also did my chores on the administrative level, as conference organizer, stints as head of my department, as president of the Dutch Mathematical Society and of the European Union Control Association. I was initiator of a graduate program (later called DISC) in Systems and Control on the national level in the Netherlands. I also did more than my share of editorial work, in particular as managing editor of SIAM Journal of Control for 5 years, and of Systems & Control Letters for the first 15 years of its existence. These professional 'service' contributions rate very high with the public. Nevertheless, I consider them as secondary to research and teaching. But I wouldn't go as far as Paul Halmos, who in his authomathography calls (his) service activities a *cop-out*.

Many personal experiences shaped my views on control, system theory, mathematics, and science, but the two dominating influences surely are: (1) my formative years as a Ph.D. student and a beginning researcher at MIT in the US, and (2) my 30 years as a mathematics professor in the Netherlands. In this article, I wish to make a few observations about the development of the field since the 1960s until now, as I experienced them personally. The research themes which I talk about are mainly those I was directly involved in.

I am not a member of what one would call the first generation of researchers in modern control theory, but of the first generation of persons who were taught a full curriculum in the area. During my years at MIT, I had the good fortune to become personally acquainted and develop close professional relationships with many of the pioneers of the subject, among

*Correspondence to: E-mail: Jan.Willems@esat.kuleuven.be

Received 30 October 2006; Accepted 8 January 2007
Recommended by S. Bittanti and M. Gevers

them were Roger Brockett, Sanjoy Mitter, the late George Zames, Rudy Kalman, Murray Wonham, Charlie Desoer, Michael Athans, and Steve Morse. With many of these people, I developed a personal friendship that lasts up to the present day.

My Ph.D. research was directed by Roger Brockett. I could not have wished for a more inspiring supervisor. I was part of a very lively and inquisitive group of Ph.D. students. We would spend weekends playing Risk or Stratego (Roberto Canales usually won), telling stories about Wiener (most of them too good to be true), and discussing the Vietnam war (we were all against it). Course work was a very substantial part of the Ph.D. program. The courses in control ranged from linear and nonlinear systems to system identification, and from optimal and process control to stochastic control with Itō calculus. They were challenging and hard work. There was also a 'minor' requirement, the equivalent of half a year's work in a department other than your own. In control, everybody chose mathematics. I took courses on abstract algebra, functional analysis including measure theory and topology, and probability and stochastic processes. It was really my first acquaintance with mathematics as a discipline, even though, as an engineering student, I had a lot of math before. I saw how you can build a rich theory starting from a few seemingly evident principles. I realized that mathematics is about concepts, not just formulas. I understood the meaning and value of rigor in proofs, and how counterexamples sharpen your mind. I finally learned what a random variable is, after having had perhaps half a dozen courses dealing with stochastics. What I mean is that I learned what it is as a mathematical notion. I realized that I had not understood the *mathematics* behind probability until then. Many years later, I realized that I had not understood the '*physics*' either.

2. Systems and Control Around 1960

In the 1960s, control was considered an electrical engineering subject, even though many applications of control involved mechanical machines or chemical processes. This could in part be explained by the fact that controllers were often implemented as electrical devices. But the mathematical methods used had a lot to do with it also.

The prevailing view of a dynamical system at that time in electrical engineering was input/output and frequency-domain based. Transfer functions were believed to be the way to characterize a system. Starting with Heaviside, symbolic calculus had been shown to be an effective tool for linear time-invariant dynamical systems. Under the influence of circuit theory, it had become evident that these methods allowed to analyze complex systems, by combining series, parallel, and feedback interconnections. The spirit of Heaviside's symbolic calculus was to be able to think of a differential operator or a delay as a formal indeterminate for which a differential operator or a delay can be substituted. Unfortunately, analysts had squeezed this marvelous idea in the mathematical *rigor (mortis)* of Laplace transforms, using complex functions, with domains of convergence and other cumbersome but largely irrelevant mathematical traps.

Electrical engineers felt more comfortable with a view of a system as a frequency transformer than with any of the equivalent time-domain descriptions. The term *filter*, referring to the fact that a system passes some frequencies more easily that others, was synonymous for system. This view was completely prevailing in control, even more so than in the neighboring areas. In circuit theory there were, after all, many nonlinear devices. Frequency-domain analysis was not especially useful when thinking about Maxwell's equations, and information theory started from an altogether different set of principles. Computers were just around the corner, but they were viewed as calculating devices. Computers as decision making machines or communication devices was something you read about in futuristic books. The black-box approach was viewed as ideal for control. Transfer functions, applied almost uniquely to continuous-time single-input/single-output systems, was the mathematical language of control. A differential equation as $p(\mathrm{d}/\mathrm{d}t)y = q(\mathrm{d}/\mathrm{d}t)u$, with p and q real polynomials, was immediately transformed to a transfer function. A number of practically very useful procedures had been developed within this framework (the PID-rules, the Nyquist criterion combined with lead/lag compensation and gain and phase margins, graphical techniques like Bode plots, Nyquist diagrams, Nichols charts, and root-locus graphs). But, by and large, there were not many surprises in all this. This framework of linear time-invariant systems left few directions for the field to grow. The absence of a sound mathematical way of thinking about polynomials and rational functions paralyzed the development of even linear multi-input/multi-output systems. From the modeling, physical, and mathematical point of view, the field was a bit narrow and a bit shallow.

In closely related domains, the transfer function point of view had met with more impressive successes. Perhaps the main problem that had been solved using transfer functions was the synthesis of passive circuits.

In his 1931 dissertation, Brune proved the remarkable result that a transfer function can be realized as the impedance of a circuit containing an interconnection of (positive) resistors, inductors, capacitors, and transformers if and only if it is rational and positive real. This result was later strengthened in a half-a-page paper by Bott and Duffin who showed that in the scalar case transformers are not needed. Later, the multivariable case was also covered, but transformerless synthesis remains an open problem in the multivariable case even today. For a while, passive circuit synthesis was applied for example to radio filter design, but nowadays solid state technology appears to have by-passed the passivity requirement as a practical concern.

Another of the great successes of the frequency-domain approach was the *Wiener filter*. The basic problem is to estimate a signal from a noise corrupted observation of it. Wiener understood that if this problem is formulated in the language of stationary stochastic processes, then this estimation comes down to constructing a filter which optimally passes frequencies in accordance to the signal-to-noise ratio. It was easy to derive this filter if for the present estimate we are permitted to use the observations for all time. But the construction of the filter which uses only the past observations turned out to be much more difficult. This is the problem Wiener solved in 1942, using mathematical methods that were perceived as unduly demanding at that time. The Wiener report, a.k.a. the 'yellow peril', because of the color of the cover and the difficulty of what was in between, was made public only after the war. It was first classified, presumably based on the principle that something that is not understood must be dangerous, a principle that intelligence agencies still honor today.

The principal systems-oriented areas in electrical engineering around 1960 were circuits, filtering, and information theory. They all had a strong presence at MIT. In circuit theory, Ernst Guillemin was the *éminence grise*. Norbert Wiener had passed away in 1964, the year before I came to MIT, but his legacy was still strong. Claude Shannon was also on the electrical engineering faculty, but he kept a low profile. I only remember one lecture by him. It was on (his) optimal investing. I was very disappointed, as I believed at that time that great scientists live from and for ideas, and do not worry about money. I was young, the words start-up and spin-off were not yet part of the daily scientific lunch-vocabulary, and the dot-com bubble was still 25 years into the future.

In the early 1960s, interest in these areas was waning. Passive circuits were becoming of secondary importance because of the transistor. Wiener did perceive a strong unity in circuits, control, communication, and information theory. In 1948, he had published *Cybernetics, or Control and Communication in the Animal and the Machine*, a 'big idea' book in which he described a theory of everything for everybody, based on systems, inputs, outputs, information flow, noise, stability, and, not to forget, FEEDBACK (in capitals). The book sold very well, and received a lot of attention from journalists, philosophers, and biologists. As a consequence, Wiener is even now often viewed as the 'inventor' of feedback. By the 1960s, many feared that Cybernetics was vacuous, and realized that it is not because something is fuzzy that it is deep. Of all these areas, information theory, with coding theory as part of it, certainly had the strongest core, both from a scientific point of view, as for its potential for applications. However, also this area suffered from overexposure and the big-idea-hype. It embarrassed those who saw its potential more realistically. Shannon himself articulated this in an editorial entitled *The Bandwagon* in the IEEE Transactions on Information Theory in 1956. The intellectual leaders in these areas felt that it was time to move on. In yet another editorial of the IEEE Transactions on Information Theory in 1958, the editor, Peter Elias, urged researchers to stop writing 'two famous papers'. One of them had a short title, *Information Theory, Photosynthesis, and Religion*. The other had a long-winded title involving a very specific application of the Wiener filter. The suggestion being that filtering was something that was understood, that there were more pressing questions to work on. Ironically, the Kalman filter, about exactly the same problem as the Wiener filter, was just about to appear. It is hard to advise researchers what (not) to do.

3. The Paradigm Shift

Around 1960, the basic model for studying dynamics in control shifted from $p(\mathrm{d}/\mathrm{dt})y = q(\mathrm{d}/\mathrm{dt})u$, to $\dot{x} = f(x, u, t), y = h(x, u, t)$. This was a major step forward. Multivariable systems could be covered without difficulty. Nonlinearities and time-variation could at least be put in evidence. Classical models from mechanics were a special case. With modest adaptations, finite state machines and automata were part of the same picture. So were, to some extent, systems described by PDE's. The input/state/output systems had much more modeling power and were far richer mathematically. By explicitly displaying its memory, the state, the model took into consideration initial conditions, something that transfer functions failed to do. The move to state space models

constituted a true paradigm shift. Control as a research activity in electrical engineering filled the void that had been left by the malaise described at the end of the previous section.

The credit for this paradigm shift must go to scientists from the Soviet Union. Perhaps because physics, mechanics, and the calculus of variations were viewed as central, or perhaps because they were used to work with differential equations, but when Pontryagin *cum suis* started thinking about control, they chose $\dot{x} = f(x, u)$ as the model for articulating optimality. They attacked the problem of trajectory planning with hard constraints on the input. The result was the maximum principle. This theorem is reasonably simple to state, but the proof is hard. The number of examples that gave nontrivial results, as time-optimal control of the harmonic oscillator, was endless. At last, there was something surprising and mathematically deep in the field.

Optimal control was picked up immediately and enthusiastically in the US, especially by mathematicians working on control, but also by theoretically oriented researchers in electrical engineering. Rudy Kalman applied the state space model to the filtering problem, basically the same problem discussed by Wiener. The results were astonishing. The solution was recursive, and the fact that the estimates could use only the past of the observations posed no difficulties. The filter gains were derived from a differential equation, which was later called the Riccati equation, as it is a generalization of a classical differential equation that went under that name. The infinite-time theory was especially subtle. The algebraic Riccati equation, a quadratic matrix equation, was the key in that case. It had multiple solutions, and sifting out the correct one involved the newly introduced notions of controllability and observability.

It was soon realized that there was a close relationship between state space models on the one hand, and transfer functions and convolutions on the other. Controllability and observability were once more key concepts. The Hankel matrix emerged as a computational tool, but this was a bit later, in 1966. The state construction problem became known as 'realization theory', and is surely one of the nicest and most important contributions of system theory. The central idea, Nerode equivalence, actually has roots in computer science. Around the same time, observers came in view, and the pole placement problem with state and output feedback was solved, with again controllability and observability as lead characters. All this came together in the linear-quadratic-gaussian (LQG) problem. The LQG controller was considered the main result in control of the 1960s. The algorithms were based on Riccati equations, the solution showed certainty equivalence and separation between estimation and control, the Kalman filter was a subsystem of the controller, and the rigorous proof required use of Itō calculus. The field was thriving.

It was in this exciting scientific climate that I did my Ph.D. research. I worked on stability problems for nonlinear systems. Roger Brockett, my Ph.D. supervisor, had invented a very clever method for constructing Lyapunov functions that yielded stability and instability criteria. Parallel to this, Irwin Sandberg and George Zames had developed functional analysis based methods, that yielded similar stability results. These were in principle more general, since they did not depend on the finite dimensionality of the underlying state space. What attracted me in this approach was not so much the generality as the transparency of the ideas. Small gains, positive operators, and conicity led to concrete results, as the circle criterion, the Popov criterion, and more multipliers than you ever wanted to see. My Ph.D. dissertation dealt with functional analysis methods for input/output stability. It was published as an MIT monograph.

That this was such a magnificent time for control is something that only became clear to me in hindsight. When one is in the middle of it, it did not seem all that special. Also, one could have expected that the professional organizations and the senior control academics were enthusiastic about what was happening. But this was not the case at all. It was a game for the young. There was a lot of skepticism. The mathematization was not welcome. Matrix exponentials were viewed as abstract mathematics (Laplace transforms were not: 'everybody' knew what they were). The fact that the circle of ideas that entered in the solution of the linear-quadratic-gaussian problem was exceptionally rich did not convince. Neither did the wealth of system theoretic concepts that had come into play. That state space systems model so many more things than transfer functions was ignored. Admittedly, there was a severe problem: the absence of concrete industrial applications. And indeed, mathematics – often mathematics for its own sake – was a more dominating motivation than technology.

4. Dissipative Systems

At that time, it was possible to be interested in all aspects of the field. The fact that the core was still relatively coherent was a big help. While trying to understand the relation between input/output and Lyapunov stability, using state construction as a guide, the notion of a dissipative system emerged. The

idea is simple. Take a state space system, $\dot{x} = f(x,u)$, $y = h(x,u)$, a real-valued function of the input and output, $s(u,y)$, called the *supply rate*, and a state function $V(x)$ called the *storage function*. If along solutions there holds $d/dt\, V(x(\cdot)) \leq s(u(\cdot),y(\cdot))$, the system is called *dissipative*. The construction of storage functions for linear systems with quadratic supply rates led to the algebraic Riccati inequality and to the acronym LMI. There were immediate relations with circuit synthesis and stochastic realization theory. Later, these ideas became important in robust control and in algorithmic methods. I was lucky. The one thing I did is to pull some things together that were in the air, to make dissipativity a concept of its own. A generalization of Lyapunov functions to open systems, to systems with inputs and outputs.

5. System Theory in Mathematics

In 1973, I moved from Electrical Engineering at MIT to Mathematics in Groningen. I spent the next 30 years there, interrupted by many short and longer visits to various institutes around the world. Such moves are never simple to evaluate, even with hindsight. But I have almost always felt fortunate to have been able to do research in a mathematics environment. The average competence level is high, there is a rich history, the subject is stable. All these factors are conducive for science. At the same time, I was never able to feel unequivocally part of the mathematics culture, where, it seems to me, too much value is put on *difficulty* as a virtue in itself. My appreciation for mathematics has more to do with its clarity of thought, its potential of sharply articulating ideas, its virtues as an unambiguous language. I am more inclined to treasure the beauty and importance of Shannon's ideas on errorless communication, algorithms as the Kalman filter or the FFT, constructs as wavelets and public key cryptography, than the heroics and virtuosity surrounding the four-color problem, Fermat's last theorem, or the Poincaré and Riemann conjectures.

One of the differences in working in a mathematics versus a theoretical engineering environment is psychological. Both fields feel insufficiently appreciated. Mathematics reacts to this by blaming the 'others', for misunderstanding them. Engineering reacts by blaming 'themselves', for not doing 'the right thing'. Discussions about theoretical engineering research often feels like visiting a graveyard in the company of Nietzsche. From the beginning of my career until now, I have always been hearing that 'the field is dead', 'circuit theory is dead', 'information theory is dead', 'coding theory is dead', 'control theory is dead', 'system theory is dead', 'linear system theory is dead', '$\mathcal{H}_\infty$ is dead'. Good science, however, is always alive. The community may not appreciate the vibrancy of good ideas, but it is there. The absence of this impatience is one of the things that makes working in a mathematics department simply more pleasant.

Mathematics has a stable core. The fashions about the level of abstraction may change, the place of statistics in the curriculum may generate disagreements, the emphasis may shift from analysis to algebra, or from the infinite to the finite, or from the continuous to the discrete. But there is general agreement about what is basic. There is also an appreciation of the historical highlights of the field, be it Banach's functional analysis, Gödel's theorem, Hilbert's list of problems, the Lebesgue integral, etc. There is something like a *communis opinio*. This is also illustrated by the widely shared interest in conjectures and open problems.

Personally, I find the difference between mathematical research and theoretical engineering more a matter of perception than substance. Take system identification, for example. In my view, it is just as much a mathematical subject, as, say, measure theory. Books or papers on system identification may use a less sophisticated and less terse mathematical language, and their authors may be employed in engineering departments, but in spirit the subjects are alike. In system identification, you want to associate a dynamical system to an observed vector time-series. In measure theory, the aim is to deal with size, volume, and integral. The intellectual scope of the former may be wider and the latter may have more applications elsewhere in mathematics, but both subjects end up by defining mathematical concepts, phrase problems using these concepts, and aim at proven results.

In The Netherlands, as in many Western European countries, the shift to state space models had not yet happened when I moved there in 1973. There was resistance to the mathematization, perceived as largely irrelevant and alienating engineers. It was a *déjà vu*, an *encore* from what I had experienced in the US years earlier. Fortunately, system and control theory was welcomed in mathematics, and by the late 1970s, all the mathematics departments of the technological universities and also of some of the classical universities had established research groups in this field. Mainstream systems and control became one of the recognized and recognizable research subfields of mathematics in The Netherlands. To some extent, this is unique in the world.

Very early on, researchers in system theory from mathematics made an effort to realize a rapprochement with control engineering departments. After

some initial hesitation, this collaboration was firmly established, thanks especially to the mechanical engineers at Delft. For over 20 years now, theorists from mathematics have been collaborating with engineering departments in a graduate systems and control program carried out on a national level. This program may not measure up, in terms of breadth, discipline of the students, and institutional commitment, to what the graduate schools in premier US universities have to offer, but it stands as an example to what can be achieved in Europe. In the mean time, there have been similar initiatives in Belgium, Italy, Bavaria, and undoubtedly elsewhere.

6. Open Versus Closed Systems

Most mathematical subjects have their origin, in one way or another, in the physical world. They derive from trying to make sense of phenomena like velocity, motion and evolution, like volume, integral and shape, like symmetry, rotation and frequency. Of course, much of mathematics has to do with unifying and generalizing all these ideas, and subsequently ask questions that come up in the new general framework. These questions often bear little relevance and completely lost reference to the original issues that led to them. In some areas, the origins remain quite clear, in others not.

One of the areas where the motivation has remained close to the surface is dynamics. From planetary motion, the n-body problem and Hamiltonian dynamics, to things like population dynamics and economic evolution, it is a seemingly small step to differential equations, $\dot{x} = f(x)$, or, more generally, to flows on manifolds. This has proven to be a very fruitful setting research-wise, leading to bifurcations, chaos, synchronization, etc.

One can, one should, ask the question if *closed* systems, as flows on manifolds and $\dot{x} = f(x)$, form a good mathematical vantage point from which to embark on the study of dynamics. In my opinion they do not, for a number of reasons. First, in a good theory the state x should be derived from a less structured model. A more serious objection is that closed systems are not good concepts to deal with modeling. A model usually consists of a number of interacting subphenomena that need to be modeled one-by-one. In these sub-models, the influence of the other subsystems needs *par force* to be viewed as external, and in principle free. Tearing leads to models that are *open*. If you view a closed system as an interconnection of two systems, these two systems will be open. Systems that take into account unmodeled external influences form therefore a much more logical starting point. Third, many basic laws in physics address open systems. For example, Newton's second law, Maxwell's equations, the gas law, and the first and second laws of thermodynamics. A good setting of dynamics should incorporate these important examples from the beginning. Finally, closed systems put one in the absurd situation that in order to model a system, one ends up having to model also the environment.

These arguments seem obvious and compelling. Twenty five years ago, it was my hope that system theory, with its emphasis on open systems, would by now have been incorporated and accepted as the new starting point for dynamical systems in mathematics. Better, more general, more natural, more apt for modeling, offering interesting new concepts as controllability, observability, dissipativity, model reduction, and with a rich, well developed, domain as linear system theory. It is disappointing that this didn't happen. What seemed like an intellectual imperative did not even begin to happen. Mathematicians and physicists invariably identify dynamical systems with closed systems.

7. Disturbance Attenuation and Model Reduction

One of the main research areas in control in the 1970s was the *geometric theory*, built around two central notions: controlled and conditionally invariant subspaces. It was an elegant theory, graceful, chic, worth dressing up for. There were also a number of seemingly convincing applications: disturbance decoupling, noninteracting control, the internal model principle, and tracking, all usually combined with stabilization. There were generalizations to nonlinear systems and to distributed systems. I introduced 'almost' versions, leading to high gain feedback. Research in this area came to a complete standstill. Geometric linear system theory is certainly the most defunct research area I have ever worked in.

The last 15 years of the twentieth century in the field of systems and control were dominated by $\mathcal{H}_\infty$-theory and by model reduction. $\mathcal{H}_\infty$ was the fourth variation of the disturbance attenuation problem to come to central stage, after LQG, or, what was later seen to be synonymous, $\mathcal{H}_2$-optimal disturbance attenuation, after bounding the effect of disturbances via pole placement and stabilization, and after exact disturbance decoupling of the geometric theory. But $\mathcal{H}_\infty$-optimal control had a number of new features that made it a very special. The solution of the

$\mathcal{H}_\infty$-optimal disturbance attenuation problem is subtle and surprising, especially because of the coupling condition between two algebraic Riccati equations. It confirmed the importance of the stability theory of the 1960s, and $\mathcal{H}_\infty$, in contrast to $\mathcal{H}_2$, leads to robustness, a central issue in control that has been neglected as a design issue. This robustness feature was reinforced by a number of new ideas on how to deal with uncertainties using linear fractional transformations, so that the small gain paradigm became applicable to a much wider class of problems. Finally, robust control brought algorithmic issues, as LMI's, convex programming, and complexity questions to the center of the field. The combination of $\mathcal{H}_\infty$, robust control, and LMI's in all their facets became a spectacularly successful research area. For a while, all these seemed like this was the only game in town. The field was thriving once again.

Late in the 1960s, it had been shown that algorithmically the problem of constructing a minimal realization from a discrete time convolution leads to rank determination and factorization of a Hankel matrix. Hard core applications-oriented people considered this problem as perhaps the most esoteric one in the field. This changed drastically in the 1980s, when it was recognized that the Hankel matrix algorithms also held the key to model reduction. The idea is to use a balanced realization, to make the system equally controllable as observable, so to speak. Combined with rank reduction using SVD and $\mathcal{H}_\infty$-bounds, this leads to effective algorithms for approaching model simplification. I believe that, because of its relevance to modeling, the bottleneck of applying mathematics, model reduction is one of the most valuable contributions of mathematical system theory to applied mathematics.

8. The Behavioral Approach

In the late 1970s, I set out to write a system theory textbook. I wanted the different topics to start on a general 'set-theoretic' level, and end up with a detailed treatment of the highly structured linear time-invariant systems. I frowned on *starting* with $\dot{x} = Ax + Bu, y = Cx + Du$, or even $\dot{x} = f(x, u), y = h(x, u)$. These were to be half-way points. Some reasoning should lead to the choice of u, y, and x. Also, I was not going to use the classical format where a definition is given first, followed by illustrative examples. I wanted this to go the other way around: show how examples lead to definitions. Of course one can fake this, by first giving the examples, expressly chosen to illustrate the definitions to come, but this is not fair: the examples had to be chosen as typical real dynamical systems.

The book was never written. The approach which I intended to follow did not work. The problem did not lay with x. State construction was quite well understood. My problem was with u and y, with the input/output partition, something that had been in the center of the field throughout the twentieth century, the very idea which Wiener saw as pulling communication and control together.

Viewing the interaction of a system with its environment in an input/output way has great intuitive appeal. It is like action and reaction. The environment acts by imposing certain variables, the inputs, on the system, and the system reacts by imposing certain variables, the outputs, on the environment. We thus arrive at the ubiquitous input/output black-box, driven by inputs on one side, and with outputs emanating on the other side. There are numerous examples and situations where this input/output view is eminently suitable, especially in signal processing, in feedback control, in adaptation, and in decision making algorithms.

But there are many awkward things with input/output thinking. We can view a typical modeling task as follows. Our aim is to model the dynamics of a number of variables. This can be visualized by means of a black-box with a number of terminals, 'places' where these variables 'live'. In many applications, these terminals are real physical entities, and there are *many* physical variables associated with *one and the same* terminal. Think of forces and torques acting on pins of a mechanical structure and displacements and attitudes of these pins. Or of currents and voltages, mass-flows and pressures, and heat-flows and temperatures. There is no sensible way in which such a modeling task leads to inputs and outputs. It demands thinking in terms of cause and effect, slippery red herrings of classical philosophy. It requires imposing a product space structure on the space of variables on a terminal. It leads to symmetry breaking. And, most importantly, it is totally inappropriate for dealing with physical interconnections.

The ubiquitous picture of a system as a black-box with an input terminal on one side and an output terminal on the other side is unfortunate for several reasons. It shows two terminals for variables that often live on one and the same terminal (voltage and current, force and position, mass-flow and pressure, heat-flow and temperature). Assume that, after having analyzed the system, we have come to the conclusion that indeed one of the terminal variables (voltage, force, pressure, or temperature) can rightfully be declared an input, and the other (current,

position, mass-, or heat-flow) an output. The classical picture suggests that the input and the output signals act at different points, whereas they act inseparably at the same physical point. It also suggests that the output signal can be directed to a different system than the one which generates the input, but this is a physical impossibility: since they are variables on one and the same terminal, they must both act between the same systems. The universal input/output black-box can pedagogically be a very misleading picture when applied to a physical system. And this becomes outright discouraging when one realizes that it is exactly the use of signal flow graphs with inputs and outputs showing the interaction pathways that is often considered by outsiders as the essence of thinking about interconnections following system theoretic ideas.

But, there is more. It happens to be a *fact*, and an obvious one when viewed in terms of units and dimensions, that physical connections identify variables of the same physical nature: voltages are identified with voltages, currents with currents, forces with forces, positions with positions, mass-flows with mass-flows, pressures with pressures, etc. So, if our physical intuition suggests that the force is an input and the position an output, then interconnection of two mechanical terminals leads to equating two inputs and equating two outputs, exactly the sort of connection that is forbidden in input/output thinking. Similarly for thermal interconnections, for hydraulic ones, etc. Input/output thinking is completely at odds with physical interconnections.

These dilemmas led to the notion of a dynamical system with simply the *behavior*, the set of legal trajectories, as the basic concept. *The behavior is all there is*. It is what a model aims at. Equivalence of models, representations of models, properties of models, approximation of models, symmetries, must all refer to the behavior. The operations allowed to bring model equations in a more convenient form are exactly those that do not change the behavior. Dynamic modeling and system identification aim at coming up with a specification of the behavior. Control comes down to restricting the behavior.

As time went on, I became only more radical in this view. Physical systems are not signal processors, and they do not interact through signal transmission. Interconnection means variable sharing, not input-to-output assignment. Control need not be sensor-output-to-actuator-input feedback. Many useful practical control devices, as dampers for vibration attenuation, heat fins, strips, and grooves to control turbulence, insulation equipment for heat or noise, stabilizers on ships, etc., do not function through sensing and actuation. Viewing the design of such appendages as part of control theory greatly enlarges the scope of our field.

As an illustration of the effectiveness of the behavior as a concept, consider the ubiquitous system theoretic notion of linearity. The behavioral definition is straightforward. A system is *linear* if the behavior is a linear subspace. Now contrast this with the definition which you find in most textbooks (and in Wikipedia). There, a linear system is defined as a linear input/output *map*. This 'map' definition does not work except in the simplest examples. Nevertheless, based on the corrupt principle that something that is flawed but presumably easy is perceived to be better pedagogy than something that is right but requires a bit of thought, this 'map' definition is almost universally used. It leads to strange conclusions. For example, with a solid dose of good-will and fast-talking, the plant $p(\mathrm{d}/\mathrm{dt})y = q(\mathrm{d}/\mathrm{dt})u$ can be construed to define a linear system in this 'map' sense. So does the feedback gain $u = -Ky$. But when this feedback is applied to the plant, we arrive at $p(\mathrm{d}/\mathrm{dt})y + Kq(\mathrm{d}/\mathrm{dt})y = 0$, and suddenly we seem to have left the realm of linear system theory by the (feed)back-door. The principal example which the map definition is motivated by is convolution, following the idea that we may as well assume that the initial conditions are zero at $-\infty$. If we apply this to Newton's second law $F = \text{mass} * a$, $a = \mathrm{d}^2q/\mathrm{d}t^2$, we obtain $q(t) = (1/\text{mass}) * \int_{-\infty}^{t}(t-t')F(t')\mathrm{d}t'$ for the relation between the position q of a pointmass and the force F acting on it. But putting the initial conditions (in this case the initial position and velocity) to zero, gets you in trouble right away. For example, when the second law is combined with the inverse square law, and with zero initial conditions, we do not get planetary motions: all planets stay put, immobilized. *Eppur non si muove*. Urbanus was right, after all.

As a last comment, it is worth mentioning the behavioral definition of controllability. A system is *controllable* if any two trajectories in the behavior are patchable, that is, if for any two trajectories in the behavior, there is a third trajectory in the behavior that has the past of the first one as its past, and the future of the second one as its sometime-future. This definition has the classical state definition as a special case. But it is simpler and more general. It generalizes immediately to PDE's. It is also more to the point. In the classical definition a system may fail to be controllable either because the control has insufficient influence on the system, or because of a bad choice of the state. With the behavioral definition this ambiguity does not occur, and controllability is independent of the representation. The system $p(\mathrm{d}/\mathrm{dt})y = q(\mathrm{d}/\mathrm{dt})u$ is controllable in the sense of behaviors if

and only if p and q have no common factors. The problem with common factors is now finally understood, it is a lack of controllability Nothing more, but nothing less. One of the urban legends of the field thus becomes a fact, a theorem. The behavioral definition of controllability has been around for about 20 years. It has essentially nowhere been adopted in teaching. It is hard to comprehend this. The cause must be sociological. It cannot be pedagogical or scientific.

9. The Physics of System Identification

Mathematical modeling is one of the most important achievements of modern science. It is at the heart of physics. As Galileo put it, *il libro della natura è scritto in caratteri matematici.* It has stimulated the development of many subfields of mathematics, physics, and engineering. It is instrumental in making computers useful devices. Mathematical modeling is the *sine qua non* of applied mathematics.

The subfield of system theory that deals most directly with mathematical modeling is *system identification.* It is a beautiful area of research. The issue is modeling from observed data. In the case of dynamical systems, the data are usually a vector time-series, and a model is sought from a chosen model class, for example, from the ARMAX systems. The 'X' refers to 'exogenous' and distinguishes system identification as it is conceived in system theory from time-series analysis in statistics and econometrics. The 'X' implies that the model class consists of open systems.

It is important to distinguish clearly between system identification and issues of system representation. There are many equivalent ways of representing a given system, of specifying the behavior. For controllable deterministic linear time-invariant differential systems, we have kernel, image, latent variable, transfer function, state space, and convolution representations. We can represent an ARMA process using its autocorrelation, its spectral density, as a function of a Markov process, and as a difference equation, a convolution or an input/state/output system driven by white noise. Passing from one representation to another, as is done in (stochastic) realization theory, is not system identification, although it can benefit greatly from understanding representation questions. A good example of this is subspace identification, which showed how state space representations can be used effectively in system identification algorithms.

System identification has grown into a mature and sophisticated area, generously supported by software and numerical toolboxes. Questions that come up in modeling from data is how to cope with uncertainty, for example in the measurements, and with the lack of fit between the data and the model. This is usually dealt with by introducing probability, by assuming that the measured data are influenced by unobserved 'disturbances', stochastic in nature.

The area of system identification is too broad to deal with in any detail here. The one item I would like to discuss briefly, is the rationale of framing this problem in the language of probability. Does probability in system identification mean *relative frequency*? Or does it mean *degree of belief*? If it means relative frequency,what is the physics which ensures that the disturbances that are responsible for the lack of fit between the data and the deterministic part of the model have a stable long run frequency? Why should there be a random mechanism at work? It is often even hard to fathom how this could come about in specific examples. If probability in system identification means degree of belief, why should we be so concerned with consistency and asymptotic efficiency? These questions can undoubtedly be given reasonable, satisfactory, pragmatic answers. What bothers me is that the textbooks on system identification do not discuss these issues at all. One cannot argue that system identification uses the wrong interpretation of probability, as the interpretation issue is usually ignored altogether. System identification is a typical area that has found comfort in mathematics, with clear questions formulated with, at best, a passing superficial reference to the underlying physics, and answers that are rigorously proven under clear assumptions.

It seems to me that the main issue in modeling from data is *approximation.* Usually we set out to model a high order, nonlinear, time-varying system, on the basis of observations that are perhaps inaccurate, sampled, and quantized. We consciously seek a model that is simple, low order, and linear. In my view, this leads to *approximation*, not *stochasticity*, as the central mathematical problem. But, before understanding approximation, we should understand exactness. So, it seems reasonable to start with identification algorithms for exact modeling, for example using the notion of the most powerful unfalsified model, then moving to (deterministic) approximation in combination with model reduction, and dealing with stochasticity at the end.

It is perhaps a bit unfair to complain about system identification for using probability without explaining why it is used, or which interpretation is used, when so many applied areas do exactly the same. Filtering

is another example. During the second world war, Wiener himself attempted to use filtering for predicting the position of enemy aircraft. I have always failed to appreciate the rationale of probability in this. I understand that enemy pilots were instructed to make unexpected maneuvers, but how does this lead, even approximately, to a stochastic process? Perhaps that is why Wiener met with little success in this endeavor. The widespread use of probability as a panacea for uncertainty, as a matter-of-fact and universal model that needs no justification nor a discussion regarding its interpretation, is for me a constant source of *Unbehagen*.

10. The Bureaucracy

Life of a researcher was simple when I came to The Netherlands in 1973. The departments left researchers in peace, the university left the departments in peace, the ministry left the universities in peace, and the EU, then called the ECC, did not dream of getting involved in science. There were no annual reports to write. The system was based on trust. But this trust was abused. The system was inefficient. Many academics interpreted the clause that research was part of their job as a friendly mild suggestion, but active researchers could concentrate on their work.

Thirty years later, these matters have completely changed. There are continuous reports, evaluations, rankings, and visitations. Networks and collaboration, top and excellent, impact, valorization, and deliverables have become buzz-words. Not only in The Netherlands, but all over Europe. The system is managed on all levels. Researchers are under great pressure to publish and compete for grants and contracts. There are, for sure, positive things that have come out of all these. Essentially everybody contributes. Idleness is frowned upon. Teaching is done with more care and thought. The number of Ph.D. students, especially in engineering, has become much larger. Research is in the public eye and is considered a necessity of modern economies.

The European Union is a case at point. The EU science directorate distributes substantial grants from contributions donated (reluctantly) by the member states. These funds are subsequently publicized as a welcome new source of research support. First you give something, then you get some of it back, and you end up being grateful. It sounds like Russell's paradox. The two main characteristics required for research programs to qualify for EU support are: *collaboration* and *excellence*. The collaboration idea is based on the belief that if you tie 2 bricks together, they will float, and if you tie 20 bricks together, they will fly. The EU aims to give its funds to *Centers of Excellence*. Centers are actually networks, called 'virtual' centers. Once a particular proposal is defined as excellent, half of the EU countries become nodes of the network, and half of the institutions in each EU country become subnodes. All researchers involved are excellent. Excellence is defined as substantially above average. For the EU, it sometimes seems like everybody can be above average.

Of course some counting, as citations, collects significant information, but only when combined with sober judgment. But many of the measures which scientists are subjected to are caricatures. The journal impact factor is a case at point. The horizon taken in its computation is so absurdly low, that with the best will in the world, one cannot take this seriously. But it is quoted over and over again, also by scientists who deal with numbers in their daily work. The publication rate is way over the top, but nevertheless the pressure to increase it is kept on. It is not unreasonable to assume that somebody who publishes two papers a year does a lot more research than someone who publishes two papers in 5 years. It is an assumption that is not without risk, but a risk one may be willing to take. Judging this sort of differences in publication rates was relevant 30 years back. But it is absurd to assume that someone who publishes ten papers a year does five times as much research than someone who publishes two papers a year. Judging such differences is what is relevant today. I know that 25 years ago, I would have defended the need for more evaluations in the educational and research system. I shiver from what it has come to.

I often wonder what the purpose really is of the enormous publication activity that goes on. Journals and conferences multiply in size and number. The work involved in preparing publications comes for a large part at the expense of time to think. In science, more writing goes together with less reading. The sheer number of publications makes it also very difficult to get acquainted with, and evaluate a new idea. I miss the emphasis on breadth and depth, on quality rather than quantity, on synthesis of ideas, on debate and scrutiny rather than passive attendance of presentations, and on reflection rather than activity. Sure, euphoria bears creativity, and skepticism paralyzes. However, questioning and criticism is an essential part of science. I have seen too many high profile areas collapse under their own weight: cybernetics, world dynamics, general systems theory, catastrophe theory, and I wonder what the future has in store for cellular automata, fractals, neural networks, complexity theory, and sync.

11. Epilogue

The field of systems and control has come a long way in the last 50 years. The mathematical methods used have expanded enormously. The techniques that have been developed for trajectory transfer, stabilization, disturbance attenuation, observers, adaptation, and robustness are deep and relevant. The modeling ideas ranging from state space systems to model reduction and uncertainty modeling are rich and versatile. The paradigm of open systems, combined with interconnection, make it into an area that fits modern technological developments well, even though systems and control has benefited less from the explosion of numerically driven and microprocessor or internet based applications than some neighboring areas, as signal processing, communication, and optimization.

We have recently seen a strong growth in the number of applications. Especially model predictive control appears to be a leading circle of ideas here. For my own taste, it has perhaps too little system theory and too much brute force computation in it, but MPC is an area where essentially all aspects of the field, from modeling to optimal control, and from observers to identification and adaptation, are in synergy with computer control and numerical mathematics.

The coherence of the field has weakened very much, and we miss a solid core. Of course, it is not an easy matter to arrive at a consensus about what the core is, or should be. The core ought to be broad. Starting teaching with classical control prepares students for a painfully limited class of problems and decision making situations. To focus on feedback control is also too narrow. Modeling and dynamics are much more pervasive scientific issues. Signal processing is for a large part based on similar concepts as system theory is. *Dynamical systems (open systems, for sure), signals, and control* is a good theme around which to build a core.

The field became enamored with mathematics. In many ways, it succumbed to the lure of looking at problems that are well defined from the start, that function in a closed environment, and where it is clear, but not necessarily easy, how to decide on correctness. In a way, similar to what drives pure mathematics. Nevertheless, at the same time, I feel that there is a strong need in the field for rigor. All too often, this is interpreted as pleading for (ϵ,δ)-type rigor or mathematical generality, but that is not what I mean. I am referring to clarity of thinking, which requires definitions, for example of a (linear) system, that truly fit, that do not introduce hidden assumptions, where a system is not identified with a representation. Rigor that translates interconnections and uncertainty from physical reality to mathematical concepts without unnecessary idealizations and unwarranted prejudices.

I have always felt uncomfortable with the distance between the mathematics of control and the physical reality which it aims at. This was extreme in the time of classical control. A transfer function is not a 'first principles' model, but at best the end result of many intermediate steps. The language of input/state/output systems is already better in this. But also this framework has been accepted without much scrutiny of how we come from a first principles physical model to an input/state/output representation.

In this essay, I have referred to 'physics' in the spirit of 'connection to reality'. It is in this sense that I feel that a convincing explanation of the physics of a problem is all too often vague or lacking, in our field in particular, and in applied mathematics in general. It is all too easy to write down a system of differential equations, argue that these equations are not unreasonable as a model for some phenomenon, subject the differential equations to a detailed analysis or a simulation which shows transition from equilibrium to periodicity to chaos, and suggest that this improves our understanding of the underlying phenomenon. Sometimes, in order to account for diversity or uncertainty, stochastics is thrown in. If the model is only vaguely justified, as it often is, for example in ecology or network dynamics, the conclusions will be even more vaguely justified. Modeling, in my view, means describing reality as literally as possible, departing from verified principles, from verified models for subsystems and their interconnections. If the basic knowledge for modeling a system is lacking, then a rough model based on guess-work has little to offer. *We do not know* is a valid answer. Unfortunately, physicists have not always given the good example, witness the speculative nature of statistical mechanics, the unanswered interpretation issues in quantum mechanics, and the nagging fear whether string theory has any connection to reality at all: it may not even be wrong, *nicht einmal falsch*.

One can articulate this paradigm for domains which, as ours, combine mathematics and physical reality, with apologies to Salieri and Strauss, as *Prima la fisica, poi la matematica*. I took this as the theme of my 'last lecture' at the University of Groningen, on January 13, 2003.

Now that we got the mathematics (more or less) right, let us get also the physics right.