The Laboratory for Information and Decision Systems

From Servo Loops to Fiber Nets

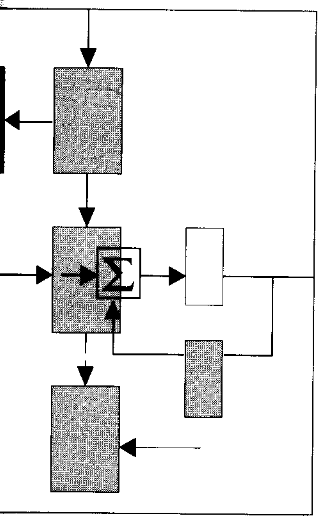

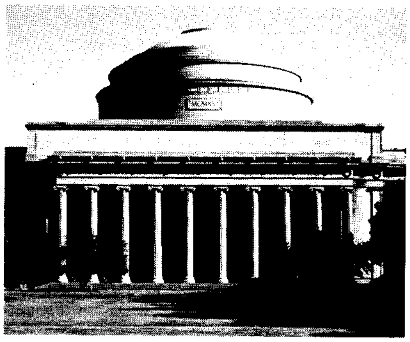

Massachusetts Institute of Technology

Preface

> ***Anniversary Celebrations*** *appear to be celebrating the past,*
>
> *but their most important function is to fix the collective identity in the present.*
>
> ***A fiftieth anniversary*** *is a story that narrates a past to support an image*
>
> *of collective identity that confirms a certain conception of the present.*[1]

The essay that follows started out as a factual account of research projects carried out in the Laboratory from its inception in 1940 to the present. It was to be a relatively short document providing some historical perspective for the fiftieth anniversary of the Laboratory. It has grown to be something more complex. It has become part history, part anecdote, and even has parts with pretensions of clarifying the structure and logic of conceptual and technological developments in one branch of engineering. Systems, Communication and Control as an engineering discipline is ultimately concerned with delineating the fundamental limitations of reliable transmission of information over a noisy channel (which today increasingly means a network of channels) and the reduction of uncertainty (arising from measurements and structure) in systems by means of feedback. The interaction of these conceptual themes with technological development is complex. It is this aspect that we have tried to understand as we have traced the development of the Laboratory in the essay. Engineering is also shaped by the history, values and culture of a society and in turn influences societal development. This aspect of engineering, especially for the fields of communication and control, is of fundamental importance but unfortunately remains untouched in this essay.

We believe that the Laboratory's contribution to both the science and technology of communication and control has been important and continues unabated today. Nowhere is this better evidenced than in the doctoral dissertations written by several hundred students while they were members of the Laboratory. Regretably we have not been able to list and give adequate credit to the many contributions made by doctoral students in the years 1940 to 1964, but we were able to be fairly exhaustive in listing doctoral students who have conducted research in the Laboratory from 1964 up to about the present. Inevitably, we have omitted mention of some people, for which we sincerely apologize. It is also our hope that we have interpreted correctly and given appropriate credit to the research performed by distinguished researchers who were formerly associated with the Laboratory as faculty and staff members.

Contributions to this essay have been made by J. Francis Reintjes, Richard S. Marcus, Douglas T. Ross (for material related to Chapter 1), Michael Athans (for material related to Chapters 2 and 3), Dimitri P. Bertsekas, Stanley B. Gershwin, John N. Tsitsiklis and Alan S. Willsky (for material related to Chapter 3), and Charles Rockland (for the section on the Nematode Project in Chapter 4). The section on Communication in Chapter 3 was contributed by Robert G. Gallager. Sanjoy K. Mitter (with essential criticism from Robert Gallager) is responsible for the sections devoted to the conceptual development of the control field and the final form of this essay. The postscript was written by Robert Gallager and Sanjoy Mitter.

This essay is being written on the occasion of the fiftieth anniversary celebration of an important laboratory of MIT. An endeavor such as this cannot be completed without the participation of people whose essential contributions often remain invisible and hence unrecognized. We want to thank Sheila Hegarty for cheerfully typing (or is it word-processing?) successive versions of this document. The "final corrections" usually turned out not to be so "final," but hopefully are "final" now. Betty Lou McClanahan helped proofread and edit the document. We thank Anne Hubbard for her patience, as we broke all records in not meeting deadlines, for her excellent composition and artwork and for her sense of aesthetics. We acknowledge the diverse contributions of Barbara Peacock-Coady over the many months leading to the anniversary celebration. Last but not least, we are grateful to Kathleen O'Sullivan for sharing the values of the faculty, staff and students of the Laboratory and believing that writing this essay was an important part of the Laboratory's fiftieth anniversary celebration. Her organizational abilities have undoubtedly been instrumental in bringing us to where we are today.

Robert G. Gallager

Sanjoy K. Mitter

---

[1] Sheldon S. Wolin, *The Presence of the Past; Essays on the State and the Constitution*, The Johns Hopkins University Press, Baltimore and London, 1989.

Co-Directors, Laboratory for Information and Decision Systems

Robert G. Gallager

Fujitsu Professor,
Department of
Electrical Engineering and
Computer Science

Co-Director, 1986-Present

Sanjoy K. Mitter

Professor,
Department of Electrical
Engineering and Computer
Science

Director, 1981-1986
Co-Director, 1986-Present

Past Directors

Gordon S. Brown

Institute Professor Emeritus,
Department of
Electrical Engineering and
Computer Science

Director, 1939-1952

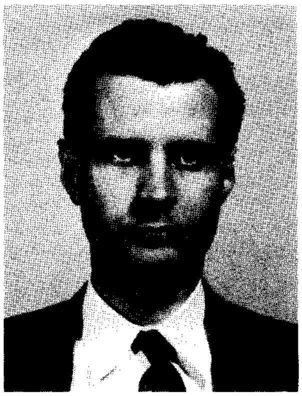

William M. Pease

Consulting Engineer,
Corporate Engineering Staff,
Raytheon Company

Director, 1952-1953

J. Francis Reintjes

Emeritus Professor,
Department of
Electrical Engineering and
Computer Science

Director, 1953-1973

Michael Athans

Professor,
Department of
Electrical Engineering and
Computer Science

Director, 1974-1981

Chapter 1

The War Years to the Sixties: From Servomechanisms to Electronic Systems

Background

The Servomechanisms Laboratory was established in 1940 by Professor Gordon Brown, the same year that the Radiation Laboratory was established at MIT and one year before the United States entered World War II. During the war years, the Laboratory was principally engaged in the application of feedback-control principles to a variety of military-related problems. Brown assembled a core research staff consisting of Jay W. Forrester, John O. Silvey, Albert C. Hall, and Tyler Marcy to work on projects such as hydraulic power transmission, servo controls for azimuth and elevation positioning of the Army's 37-mm gun mount, the design and construction of a fusesetter rammer, and the design and construction of mount power drives for 40-mm guns. They were assisted by several graduate students and by four naval lieutenants assigned to the electrical engineering department for a special program in servomechanisms and fire control. All four lieutenants — Edwin B. Hooper, Lloyd Mustin, Horacio Rivero, and Alfred G. Ward — did their theses under Brown's direction and made important contributions to the Laboratory's programs through their thesis research.[1]

Systems, Control and Computation Research

In the period immediately preceding 1940, perhaps the three papers most influential in laying the conceptual foundations of Feedback Control Systems were Minorsky's paper "Directional Stability of Automatically Steered Bodies," Nyquist's paper on "Regeneration Theory" and Hazen's paper "Theory of Servomechanisms." Thus the term and even the concept of Servomechanisms was very much in the air in 1940 when the Servomechanisms Laboratory was founded. Indeed, formal classroom instruction in control had begun in the Electrical Engineering Department in 1939 (prior to the founding of the Servomechanisms Laboratory) as a two-term graduate sequence entitled "Theory and Applications of Servomechanisms." This subject was taught for many years by Professors Gordon Brown, Donald Campbell, William Pease, George Newton and Leonard Gould. An additional undergraduate elective subject, first entitled "Industrial Applications of Servomechanisms" and later called "Feedback Control Principles," was presented for the first time by Campbell in 1947. Five years later Campbell introduced the subject in "Process Control", which was followed some six years later by "Chemical Process Control" taught by Gould. Research in Servomechanisms performed in the Laboratory and the experience of teaching the subject led to its codification in the form of a textbook by Gordon Brown and Donald Campbell entitled *Principles of Servomechanisms* in 1948.

Related intellectual developments were taking place at about the same time, most notably: the work of H.W. Bode on Feedback Amplifier Design (at the Bell Telephone Laboratories); the work on feedback control in the Radiation Laboratory, as outlined in *Theory of Servomechanisms* by James, Nichols and Phillips; and the pioneering work of Wiener and Kolmogoroff on filtering and prediction of stationary time series. Even the essential unity of the fields of communication and control, a theme which the Laboratory would champion in the seventies was no secret to the masters. It is interesting to quote from Wiener's book on Cybernetics

> *On the communication engineering plane, it had already become clear to Mr. Bigelow and myself that the problems of control engineering and of communication engineering were inseparable, and that they centered not around the technique of electrical engineering but around the much more fundamental notion of the message, whether this should be transmitted by electrical, mechanical or nervous means. The message is a discrete or continuous sequence of measurable events distributed in time—precisely what is called a time series by the statisticians. The prediction of the future of a message is done by some sort of*

---

[1] For a historical account of the Servomechanisms Laboratory, see *A Century of Electrical Engineering and Computer Science at MIT, 1882-1982*, Karl L. Wildes and Nils A. Lindgren, The MIT Press, Cambridge, Mass., 1985.

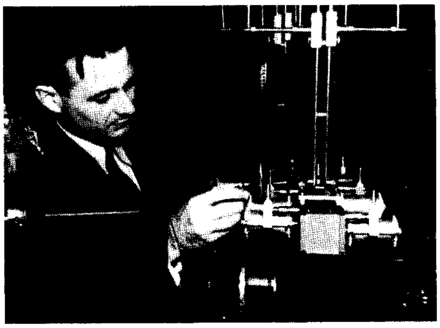

Gordon S. Brown

George C. Newton

> *operator on its past, whether this operator is realized by a scheme of mathematical computation, or by a mechanical or electrical apparatus.*[1]

Brown and Campbell's book on Servomechanisms uses a linear differential equation description of dynamics which is then transformed into the frequency domain by means of the Laplace Transform. The frequency domain analysis and synthesis as developed in the doctoral thesis of A. C. Hall, written in 1943 and entitled "The Analysis and Synthesis of Linear Servomechanisms," was incorporated in this book. The approach to Servomechanisms design in the book might be called the "trial and error method" and is primarily concerned with questions of stability and transient response of feedback systems.

Nine years after the publication of Brown and Campbell's book, Newton, Kaiser and Gould's *Analytical Design of Linear Feedback Controls* was published. This book is in some sense a report on the research then being carried out in the Laboratory. In this book, the revolutionary step of combining the Kolmogoroff-Wiener theory of filtering and prediction of stationary time series and the theory of servomechanisms was undertaken with a view to understanding the fundamental limitations of linear feedback control systems. The contents of the book and the objectives of writing it can best be described by quoting from the authors' preface:

> *According to analytical design theory, the best compensation for a feedback control system is implicitly determined by the specification. When the designer applies the analytical design method he proceeds directly from the problem specifications to the compensation that minimizes or maximizes the specified performance index. By this method the design is accomplished once and for all without recourse to a series of trial-and-error designs.... Unfortunately, computational effort is often greater with this method than with the trial-and-error design procedure. Thus we frequently find the most effective approach to be a combination of the two techniques.* [2]

Newton, Kaiser and Gould go on to say that they are interested in understanding the factors that fundamentally limit the performance of systems—factors such as input noise, disturbances, non-minimum-phase fixed elements, and saturation in the fixed elements and that such an insight cannot be had through the study of the trial-and-error design procedure.

It is interesting to note that in the last few years we have returned to the fundamental issues discussed in Newton, Kaiser and Gould's book, albeit, in the guise of H^2, H^∞ and L^1 optimization. It would be important to pay attention to their cautionary note about the use of this optimization theory.

> *Control systems often employ mechanical, hydraulic, or pneumatic elements which have less reproducible behavior than high quality electric circuit elements. This practical problem often causes the control designer to stop short of an optimum design because he knows full well that the parameters of the physical system may deviate considerably from the data on which he bases his design.*[3]

Hardware and Software Applications

We have so far been highlighting the theoretical and conceptual contribution of the Laboratory in the domain of feedback control systems from its inception up to about 1960. These developments took place initially from the needs of World War II in gun control, radar systems and communications and in the latter part of this period, from the needs of the U.S. space program. Its mathematical foundations lie in complex function theory and harmonic analysis. Its creativity lies in the discovery of the hidden conceptual structures behind engineering problems and in

---

[1] Wiener, N., *Cybernetics or Control and Communication in the Animal and the Machine*, The MIT Press, Cambridge, Mass., 1948, pp. 8-9.

[2] Newton, G.C., Gould, L.A. and Kaiser, J.F., *Analytical Design of Linear Feedback Controls*, John Wiley and Sons, 1957, pp. 6.

[3] Ibid, pp. 23.

Leonard A. Gould

Jay W. Forrester

crystallizing them through the introduction of appropriate mathematical structures. But the interaction between theoretical and conceptual ideas, engineering synthesis and technological development in the field of systems, communication and control is more complex. It is in fact a highly complicated feedback process. Conceptual developments in engineering are incomplete until they lead to a new algorithm, new apparatus or machine. These in turn require new conceptual ideas for their full utilization. This complex interaction is best exemplified by work in the Laboratory on the Brookhaven project, Whirlwind project, the Numerical Control project and the Computer-Aided Design project.

The Brookhaven Project

After the war the Laboratory continued its research on military-related problems, but new opportunities for peacetime applications of automatic control appeared. Two major projects, the Brookhaven nuclear reactor project and the Whirlwind computer project, constituted the major part of the Laboratory's efforts in the middle and late 1940's. The Brookhaven National Laboratory, which had responsibility for developing the first peacetime nuclear reactor for scientific research purposes, turned to the Servomechanisms Laboratory for the design and building of the power drives and controls for actuating the reactor rods and for associated instrumentation. The control-system work was headed by Pease (the second director of the Laboratory), and the instrumentation was under the leadership of Professor Truman S. Gray. This was a major undertaking, amounting to an expenditure of some $1 million and 185,000 pounds of equipment. Much to Brown's satisfaction and to the credit of all who contributed to the project, when the reactor was brought up to critical state, it performed as expected. Upon completion of the Brookhaven Project, Professor Gray continued his research in electronics for nuclear particle detection in collaboration with Professor Albert Van Rennes.

Whirlwind Project[1]

In late 1944 the Laboratory responded to the need of the Navy for an electronic computer that could simulate the dynamics of aircraft. It was envisioned that such a computer would be a useful device for training pilots and an analytical tool for designers of new aircraft. The project, which was called the Aircraft Stability and Control Analyzer project, was headed by Jay Forrester, now Germeshausen Professor Emeritus of Management. The Analyzer was originally conceived as an analog machine, but during the summer of 1945 thinking shifted away from the analog domain and toward digital techniques as a means of overcoming formidable obstacles being encountered in implementing the analog computations. Within two years, Forrester was able to demonstrate his first digital machine, which he dubbed WHIRLWIND I. An advanced version of the machine was developed at the Barta building on Mass. Ave., a block or so from the MIT museum. Interestingly, WHIRLWIND was never used as an aircraft simulator or cockpit trainer but played a key role in the development of the nation's continental air defence system.

Project Whirlwind contributed to computer technology in an important way but also epitomized the close connection that exists between developments in computers and computation and the science and technology of control systems. The most famous contribution of the Whirlwind project was the random access, magnetic core storage feature of the Whirlwind machine. From our perspective, no less important was the idea of incorporating a computer in the feedback loop.

Numerically Controlled Milling Machine Project

With control theory and applications projects and the Whirlwind digital computer project existing side-by-side in the Laboratory during the 1940's it was inevitable that the disciplines of control and digital computation would converge sooner or later in a single application. This happened with the initiation of the Numerically Controlled

---

[1] For a detailed study of Whirlwind, see *Project Whirlwind: The History of a Pioneer Computer*, Kent C. Redmond and Thomas M. Smith, Digital Press, 1980.

Douglas T. Ross, John E. Ward

Milling Machine project in 1949. Although the project goal was sharply focussed on the development of a digitally controlled, three-axis milling machine, the results had far-reaching implications. The successful demonstration of the MIT Numerically Controlled Milling Machine marked the beginning of the digital computer as a control element in feedback control systems. Applications of real-time digital control eventually spread into such diverse fields as chemical process control, space guidance and control systems, computer-controlled flexible manufacturing systems, and a host of other applications where real-time control is required.

The Numerical Control project, which was under the leadership of William M. Pease and later James O. McDonough, culminated in 1952 in an experimental model of a three-axis continuous-path milling machine that could sculpture mechanical parts at conventional cutting speeds with minimal human intervention. This milestone had a major influence on the future of the machine-tool industry, for it put the industry onto an entirely new growth curve. Henceforth the industry would be closely intertwined with the electronics industry and would grow in accordance with the growth and progress in electronics. A patent was awarded to Forrester, McDonough, Pease, and Susskind for the experimental machine tool.

The APT System and the Beginning of CAD

An economic study of numerical control for metal-removal purposes conducted in 1953 by Professor Robert H. Gregory, a Sloan School faculty member, and his colleague Thomas V. Atwater, Jr., concluded that, because programming for numerically controlled machine tools was a costly item, the technology would be more readily embraced if more cost-effective programming could be found. Follow-on research in automatic programming for continuous-path machine tools was therefore undertaken in 1956 with Douglas T. Ross, head of the Computer Application Group, as principal scientist and John E. Ward as project administrator. Emerging from this effort was the APT System of programming, APT standing for Automatically Programmed Tool system. APT was a special purpose programming language—among the first, if not the first elaborate programming system to achieve an excellent match between laymen-users of computers and the programs they were using. Machine-tool commands were expressed in English-like words, thus making it easy for persons with no programming skills to engage the system. All complex communications and computations were carried out automatically at a high-level language stage without the need for significant computer expertise at the user level.

In response to an urgent Air Force preparedness need, starting in mid-1957 and using the Whirlwind computer work as a continuing research base, the APT Project designed and led a nationwide effort, in collaboration with the Aircraft Industries Association, in which programmers (55, overall) at some 19 companies wrote subsystem components for an APT System that could be used to put newly-acquired commercial machine tools into productive use. The work was coordinated from MIT, where the Control Element component also was written, and the various subsystems were assembled, corrected, and tested on the IBM 704 computer at MIT, for distribution back to the participating companies. The first release was a 2D system in the summer of 1959. APT was an "open system" in today's parlance, with a standard output that could drive any controller/machine—tool combination for which a postprocessor has been written. In 1960 the AIA gathered programmers in San Diego to continue the work as the AIA APT Project, with MIT providing a new, comprehensive, guaranteed-within-tolerance 3D tool-control calculation capability, and allowing arbitrary-shaped cutters.

Even at the beginning of the APT Project, it was clear that sculpturing mechanical parts by means of numerically controlled cutting tools was only part of a larger process that began with the specification and design of the part itself. The next logical step, therefore, was to get the computer involved in the overall manufacturing process at the earliest stage possible. Led by Ross, the Computer-Aided Design Project began in 1960 with a five-year collaboration with the Mechanical Engineering Department's Design Division, led by Professors Robert

W. Mann and Steven A. Coons, which overlapped with further joint work (1963-1969) with Project MAC (MIT's pioneering time-shared computing project, led by Prof. Robert M. Fano). Emphasis was placed on software development as well as high-performance computer displays, work led by John E. Ward, who had joined the Lab in 1946, and who had had a long string of practical accomplishments in the control/computer field. The first general-purpose, machine-independent software engineering programming language (object-oriented, with pointers, n-component data elements, and integrated packages of callable routines), AED (Automated Engineering Design or Algol Extended for Design — pronounced "aid") and its integrated libraries and tools for generating application-specific languages and systems came from Ross's Computer Applications Group. The first interactive 3D display system came from Ward's Display group, from the thesis of Robert M. Stotz. Coons and Mann laid out and popularized the principles of CAD as practiced today, and the "Coons Patch" pioneered parameterized shape description. Industry collaboration also continued and expanded. From 1964 to 1969, some 30 programmers from 21 organizations worldwide contributed 362 man-months of effort as visiting staff of the AED Project, through the AED Cooperation Program, and many major MIT projects using MAC facilities did their work using the AED software technology (software-engineering language and discipline, automated system-building tools, and libraries of portable and reusable software components)—many using the ESL Display Console (dubbed "the Kludge") as an attached processor to the MAC system. The Kludge was probably the first graphic display capable of being operated on a time-shared computer. Notable areas of application were circuit simulation by AEDNET (Katzenelson) and CIRCAL (Dertouzos), dynamic system simulation by DYNAMO (Forrester), econometric modeling by TROLL (Kuh), ship design (Hamilton and Weiss), and 3D molecular structure of proteins (Levinthay), etc. Ross supervised the first master's theses related to mechanical CAD (A.F. Smith, Mathematics Department) and electronic CAD (Meyers) in 1960, and taught the first graduate course in software engineering (SofTech). AED underwent several releases to industry on the IBM 7094 , Univac 1108, IBM 360 systems, with cross-compiling to DEC PDP-10, GE Multics, and others.We take for granted today that engineers will have workstations sitting on their desks, will have access to sophisticated graphics, and will program in a high level language particularly attuned to the problem domain of interest to the engineer. It is startling to see the presence of all these ideas in CAD work in the Laboratory carried out thirty years ago.

Toward Electronic Systems

There were other research projects in the Laboratory which represent the interaction between theory and practice, conceptual ideas and technological development, which we have referred to before.

Project Porcupine

The Porcupine project, a project of the 1950's, was an outgrowth of Lincoln Laboratory's charter to engage in research related to air defense. Porcupine was conceived as a low-cost, short-range, point-defense system for use against low-flying aircraft. Simplicity was supposed to be its strong point. Its major components were a Pulse-Doppler radar for detecting and tracking aircraft down to altitudes of fifty feet or less, a fire-control computer (which turned out to be an analog computer), and an egg-crate-like bundle of short-range rockets that flew unguided to their target. The project was sponsored by Lincoln Laboratory, and Lincoln developed the Pulse-Doppler radar under the leadership of Reintjes and Louis Smullin who, at the time was a Lincoln Lab staff member. Servo Lab staff members, including Jack Simons, Mark Connelly, Michael Fitzmorris, and others, did the initial system studies and were responsible for the fire-control computer; Lincoln Lab procured the rockets and launcher from an outside source. Porcupine was a big system development with at least one notable achievement: The Pulse-Doppler radar was the first of its kind that could automatically detect and track low-flying aircraft in the presence of heavy ground clutter, and this capability was dramatically demonstrated in flight test. The equipment was transported to the Naval test station on the Mojave Desert where it knocked down two drones flying at 50-foot altitude in two tries. A weapon system with a demonstrated hit probability of 1.0!

The Radar Project

When Professor Reintjes assumed leadership of the Laboratory in 1953 he initiated a research project in radar in order to diversify the Lab's activities and to satisfy his own research interest in this area. Joining him in this endeavor was Godfrey Coate, his Radar School colleague and coauthor of their textbook on principles of radar. In addition to Coate the project was under the leadership of Richard Spencer, Lawrence Swain and Jack Silvey. At the outset, the project goals were to conduct systems investigations leading to sound foundations for the design of advanced radars for aerospace vehicles, and to devise and investigate methods by which modern materials, devices and techniques could be used advantageously to enhance the performance and useful life of radars and to reduce their size and weight. The establishment of all-solid-state radars as the norm was the ultimate objective. Heavy emphasis was placed on high-power pulse modulator research, since this was the subsystem where the bulk was and very little had been done to reduce it. The cooperation of industrial suppliers of solid-state

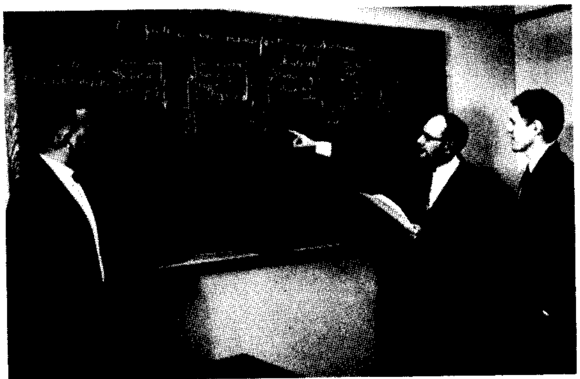

J. Francis Reintjes, George C. Newton, and John E. Ward

components was sought and received; as results were attained, radar manufacturers from industry were invited to assign representatives to join the group so that the new technology could be transferred into industrial settings. Solid-state technologies for spacecraft radars were also investigated under the auspices of NASA, and extensive flight testing of an experimental model was conducted over a large area of western United States under the leadership of James R. Sandison. It was on this project that James K. Roberge performed his doctoral research which led to his textbook *Operational Amplifiers: Theory and Practice.*

As industry assumed responsibility for solid-state radar developments, the research of the group turned to analytical studies of advanced radar system techniques. In particular, heavy concentration was placed on digital signal processing for side-looking airborne radars for high-resolution ground mapping. Over a span of twenty years the project provided research opportunities for some 50 thesis students at the undergraduate and graduate levels.

Hybrid Analog-Digital Computation Techniques

Under the leadership of Mark E. Connelly the Laboratory undertook a series of investigations in the realm of hybrid computers and computation in the 1960's. The overall objective was to exploit the relative simplicity of this computational tool over general-purpose digital machines for specific applications such as, for example, simulation of the dynamics of aircraft for operational flight trainers. The project efforts culminated in versatile Laboratory facilities that included a family of high-speed pulsed-analog elements capable of performing a wide range of computational operations.

By this time it was clear that the Laboratory was conducting research on a front which was much broader than its original charter, especially if servomechanisms is interpreted in a narrow sense. In 1959 the Laboratory had acquired a new name, *The Electronic Systems Laboratory,* to better reflect this reality. In the next section we describe an important project which clearly demonstrates this broadening of research activities of the Laboratory. This research is continuing even today, but it is placed here so as not to break the historical continuity of the research theme. It is also placed here because the style of doing research in this project is very much that of this early period of the Laboratory's existence.

Information Transfer and Information Processing

The Laboratory's research programs in this area began in the mid 1960's with Project Intrex. Intrex, standing for INformation TRansfer EXperiments, was the brainchild of Carl Overhage, Professor of Engineering and former director of Lincoln Laboratory. The goal of the program was to exploit the digital computer and other digital and analog equipment as a means for increasing the effectiveness of the library as an information transfer facility. The Intrex investigations can be conceived as deriving from, and extending, certain cybernetic themes from earlier Laboratory work as well as initiating or supporting a number of emerging trends that have become increasingly paramount in the current computer age. The Laboratory undertook technical responsibility for the research program under the leadership of Professor Reintjes, assisted by Richard S. Marcus. The research centered on two topics: the development of a computer-stored augmented catalog of information on the technical journal literature, and the development of a microfiche-stored set of articles which corresponded to those in the catalog and which could be retrieved and displayed remotely on video-display tubes. Documents could also be given to the users in hard-copy form. The experimental subjects were users who had a bonafide need for the catalog and full-text information stored in the database of the Intrex system. Intrex was unique in that it addressed the problem of retrieving easily and quickly information being sought in professional journal articles, a class of literature not indexed in depth in library card catalogs.

The inherent ambiguity in searching in natural language texts is associated with the fact that in document retrieval, in contrast with lookup in numerical tables, there is generally no single "right" answer, and the question itself tends to evolve as feedback is obtained. This leads to the conception of the computerized information retrieval

Richard S. Marcus

process as essentially an interactive one in which the man and the machine must have intimate cooperation in order to achieve maximum effectiveness and efficiency. Sometimes the human will direct operations and sometimes the computer will recommend, or even enforce, decisions — as well as carry out mundane processing tasks — in the optimum "mixed-initiative" modality.

In this framework the Intrex experimental system was designed so that "ordinary" (i.e., even computer-inexperienced) users could operate it easily and effectively. Thus a simplified command language, extensive computer-provided instruction, and specially designed terminals for catalog and text reception mirrored earlier lab work on the Kludge and Ross' APT language and helped usher in the period of "user friendly" systems and the push toward multimedia computer environments. The Intrex model of a fully automated library with guaranteed, simultaneous remote access by multiple users is not yet completely achieved, but still serves as a beacon for future development. Intrex' demonstration of electronic storage and transmission of full document text (including graphic and photographic — cf bit-mapped — images), together with later work by Reintjes in the digital encoding and transmission of document images, meshed with the facsimile developments of that time and presaged the modern world where "fax it" has become the buzz word.

The Newspaper Project, spanning the ten-year period 1966-1976 and sponsored by the American Newspaper Publishers Association, was one outgrowth of the Laboratory's experience with Intrex. Like Intrex, which intensified the interest of the library community in computers, the newspaper project served to convince publishers and their editorial staff of the merit of computer-aided processing of the news and the production staff of the role the computer could play in the manufacturing of their newspapers. Just as computers are extensively used in library operations, so also are they used in a wide variety of newspaper operations, such as composing, editorial operations, classified-ad taking, typesetting and business operations. The project served as a catalyst to bring the newspaper business into the computer age and provided thesis opportunities for many students, two of whom chose the newspaper business for their careers.

Another outgrowth of the Intrex Project has been the continuing effort by Marcus to make document retrieval procedures more rational, effective, and comprehensive, as well as simpler for the searcher. One aspect of these investigations addressed the problem of accessing the world's literature which is to a considerable extent now referenced in hundreds of computerized bibliographic databases but whose access is limited by the multiplicity, dispersion, heterogeneity, and complexity of the dozens of retrieval systems which maintain these databases. The answer to this problem has been sought in the conception of a "computer intermediary" system that would "talk" to searchers in a common, easily-mastered interface language, help them identify databases of potential interest, translate their informational requests into the appropriate command language of a system maintaining one of those databases, automatically connect to that system and transmit the commands, and report back the results in an understandable format to the searcher. The intermediary system would then help the searcher modify his/her request in order to obtain more relevant results. Several experimental intermediary systems, under the generic name "CONIT", were built and in one series of controlled experiments it was demonstrated that retrieval-system inexperienced searchers using CONIT were able, on average, to obtain as many relevant references from on-line databases as they could when working with human expert information specialist intermediaries on the same topics with the same databases. CONIT as an expert computer assistant paralleled, in certain respects, the goals and techniques of the so-called "expert system" field of the artificial intelligence discipline.

The more recent and current work on the CONIT system has sought to investigate the further elaboration of "intelligent" techniques for the document retrieval application so as to go beyond the straightforward simulation and simple matching of what human experts can now accomplish. Achievement of this goal is being pursued by developing a quantitative model of the factors influencing the search process and incorporating that model into a new assistance system where the computational prowess of the computer is integrated with the human searcher's intelligence. Some capabilities of this new retrieval assistant, unique to Boolean-based systems, include automatic ranking of documents by their estimated relevance to the topic; dynamic, quantitative evaluation of the search effectiveness (including an estimate of the number of relevant documents not yet found); and an optimized search strategy reformulation based on minimal searcher relevance feedback information. An advanced CONIT system incorporating these features in a computer workstation environment is now being completed in preparation for further testing.

Chapter 2

The Sixties to the Early Seventies: The Golden Age of Control

The New Intellectual Climate

Towards the end of the fifties, external events occurring in the world and the internal crisis in the conceptual foundations of control theory were to generate forces that would dramatically change the course of control research in the sixties. Sputnik 1 was launched by the USSR in 1957 and the US Space Program was initiated shortly thereafter. Solutions were needed for new control problems, such as the attitude control of satellites and guidance and control of space vehicles. The natural description of these control problems was in terms of systems of linear and non-linear ordinary differential equations in the "state" variables representing the equations of motion of bodies. The linear theory of the fifties, with its focus on input-output descriptions of time-invariant systems and frequency domain formulations, was no longer adequate to handle these new control problems. Problems of trajectory determination of space vehicles naturally led to the formulation of control problems via an optimization criterion in the time-domain. The new field of Optimal Control was born. There was renewed interest in the Calculus of Variations, an old subject, but extensively developed in the USA in the thirties by the Bliss School at the University of Chicago. The new issue in optimal control, perhaps not sufficiently emphasized in earlier Calculus of Variations theory, was the presence of inequality constraints on the control and possibly state variables. An early example of this can be found in the work on Time-Optimal Control by Bushaw and later by Lasalle. In the Soviet Union, the theory of optimal control was taken up by Pontryagin and his co-workers, this led to the celebrated Maximum Principle as a necessary condition of optimality. In the United States, Bellman was to develop Dynamic Programming, a general method for treating both deterministic and stochastic optimization problems such as those arising from control problems. Its roots can be found in the Hamilton-Jacoby-Caratheodory view of classical mechanics (a fact noted by Kalman) and Wald's Sequential Analysis. Thus, by the early sixties, optimal control received an almost definitive treatment in the hands of Bellman, Pontryagin and others.

We now describe some of the symptoms of the internal crisis in the field. Feedback control systems, designed using existing theory, often gave rise to hidden instabilities that the theory did not predict. There was no satisfactory theory for dealing with multivariable control systems. Attempts to solve these problems via decoupling, thereby reducing them to single-input single-output systems, were by and large failures (indeed the decoupling problem would only be solved in the seventies in successive stages by Falb and Wolovich, Gilbert, and Wonham and Morse). Finally, the solution of the Wiener filtering problem could only be effectively carried out for random signals with a stationary spectral density and even then this method of solution was computationally unattractive. The solution to these problems would require a conceptual breakthrough. This was provided first by introducing the idea of an internal dynamical description of a control system via ordinary differential equations describing the evolution of the state variables of the system and second by introducing the theory of Markov Processes (the probabilistic counterpart of the notion of state in a deterministic context) to solve the filtering problem. Indeed, the concept of a solution as an algorithm was an essential contribution of this time. Moreover, there was a reconciliation of the input-output and state space points of view. This required the introduction of new conceptual ideas of controllability, observability and minimal internal realizations of an input-output map. The early paper of Gilbert on controllability and observability, the seminal work of Kalman on controllability, observability and realization theory, the work of Youla on realization theory (to cite some key contributions) served to remove the crisis atmosphere of the late fifties.

Kalman and Bucy solved the Wiener filtering problem by restricting attention to random signals with a rational spectra and modelling them as functions of a Gauss-Markov process. At the same time, they considered the more general non-stationary situation. The Gauss-Markov process is described by a linear differential equation forced by white noise. The resulting filter was itself a linear differential equation fed by the so-called innovations

Roger W. Brockett

George Zames

process. The time-varying gain of the filter was given in terms of the solution of a non-linear differential equation, the celebrated Riccati equation. Kalman also studied the asymptotic behavior of the filter, and the equilibrium solution solved the Wiener filtering problem. It was perhaps inevitable that the two strands - optimal control and optimal filtering would come together. This came about by considering a special case of the optimal control problem, the so-called linear quadratic (LQ) problem and its stochastic version, the linear-quadratic-gaussion (LQG) problem. A discrete-time version of the last-mentioned problem was investigated by Florentin in England and Joseph and Tou in the U.S.A. They showed that the solution of the LQG problem separated into a filtering part (which can be solved via the Kalman-Bucy filter) and a control part (which can be solved by LQ theory). A rigorous definitive treatment of the LQG problem which used the full machinery of stochastic calculus was given by Wonham somewhat later in the sixties.

Contributions of the Laboratory in Perspective

To give a sense of the Laboratory's early presence in the developments sketched above, we cite the pioneering theses of C.W. Merriam and W. Kipiniak. Merriam's thesis, written in 1958, was entitled "Synthesis of Adaptive Controls". The single-input Quadratic Control Problem was formulated in the thesis and solved using the method of dynamic programming. The celebrated Riccati equation of quadratic control appears in this thesis (probably for the first time). Merriam continued this research in the General Electric Company and subsequently published his book *Optimization Theory and the Design of Feedback Control Systems* in 1966. Kipiniak's thesis entitled "A Variational Approach to Dynamic Optimization for Control System Synthesis" was completed in 1960 and published in book form as *Dynamic Optimization and Control* in 1961. In contrast to Merriam's work which used the method of dynamic programming, Kipiniak used the theory of calculus of variations in order to obtain necessary conditions for quite general optimal control problems (including problems with time-delay and distributed parameter systems) and also considered approaches to the numerical solution of the resulting two-point problem. The last part of Kipiniak's thesis is a valiant attempt towards formulating and solving a non-linear stochastic control problem.

In the early sixties, a very significant boost to the research of the Laboratory in the field of so-called Modern Control Theory took place when Michael Athans, Roger W. Brockett, and George Zames were appointed as Assistant Professors in the Electrical Engineering Department with primary affiliation with the Electronic Systems Laboratory. Michael Dertouzos was also appointed as a faculty member in the Department and had his primary affiliation with the Laboratory.

Michael Athans, after receiving his Ph.D. from the University of California at Berkeley in 1961, joined the technical staff of the Lincoln Laboratory in Lexington, Mass. At Lincoln he worked under the supervision of (the late) Fred C. Schweppe — (who joined the MIT faculty later) — in a small group that included now well known researchers such as Harold J. Kushner, Peter L. Falb, Leland Gardner, Joel Moses, Tom Bartee, John Lewis, and Harold Knudsen. Athans' work at Lincoln revolved arround optimal control and estimation problems. He co-authored the classic text *Optimal Control* with P.L. Falb, and he gave an informal seminar course on optimal control at MIT during the preparation of the book. At the invitation of Professor G.C. Newton he joined the faculty in 1964. At about the same time, Roger W. Brockett, after completing his Ph.D. from the Case Institute of Technology, under the supervision of M. Mesarovic, also joined the MIT faculty. Together with George Zames, they were assigned the development of a new graduate level curriculum which reflected the recent advances in modern control theory and the strengthening of the graduate research area in systems and control.

The sixties were indeed a period of great intellectual excitement for the control group at the Electronic Systems Laboratory. From an academic perspective, Roger W. Brockett took responsibility for developing the first course

Michael Athans, Sanjoy K. Mitter, Fred Schweppe, and Jan C. Willems

on linear systems. His important book *Finite Dimensional Linear Systems* was based upon the material taught in that course. Roger Brockett and George Zames taught courses on stability theory and non-linear systems. Michael Athans developed a course on optimal control. Leonard Gould continued to attract students from both Electrical Engineering and Chemical Engineering in his process control course. The book *Chemical Process Control: Theory and Applications* is an outgrowth of this course. These core subjects, followed by more advanced subjects, attracted a large number of first-rate students who did their dissertations in this area. Some of the key doctoral research during that period was carried out by P.R. Belanger, J. Burchfiel, R. Canales, A.S. Debs, T.L. Fortmann, D.L. Gray, S.G. Greenberg, M. Gruber, D.L. Kleinman, H. Kurihara, W.S. Levine, A.H. Levis, M. Muraay-Losso, J.B. Plant, G. Prado, A. Rahimi, G. Skelton, R. Skoog, J.C. Willems, J.L. Willems, H.S. Witsenhausen - to mention just a few. Some of the research accomplishments of these students were: the control of uncertain multi-agent systems (Witsenhausen), the solution of the so-called output feedback problem (Levine), the development of computational methods for time and fuel-optimal control (Plant, Gray), Kleinman's method for solving the algebraic Riccati equation using Newton-Raphson type methods, the counterintuitive properties of optimal multivariable process controllers (Kurihara), optimal control of sampled data systems (Levis), non-linear stability theory (the Willems brothers, Gruber), time-varying network theory using state-space methods (Skoog).

In addition to the intellectual fervor generated by the students, a large number of distinguished visiting faculty contributed to the intellectual excitement, such as L.W. Polak (from the University of California at Berkeley), H. Rosenbrock (from the Manchester Institute of Science and Technology), and H.J. Kushner and W.M. Wonham (from Brown University). A series of monthly meetings attended by faculty and graduate students from MIT, Harvard and Brown were initiated. We all benefited from the —often caustic— remarks by Professor Solomon Lefschetz of Brown University.

During the same period, research in the Electronic Systems Laboratory consolidated numerous theoretical results in linear and non-linear optimal control, filtering, and related computational research (including CAD software). Applications-oriented studies continued in the area of chemical process control, satellite control, and automatic control of high speed trains under the auspices of Project Transport. The Laboratory also did pioneering research in the area of non-linear systems and stability theory. George Zames wrote his doctoral dissertation in the MIT Research Laboratory of Electronics on "Non-linear Operators for System Analysis". In this work Zames adopted an abstract functional-analytic view for studying non-linear systems. The seeds of input-output stability theory, developed by him while he was at the Electronic Systems Laboratory (1961-1965) and later at the NASA Control Theory Center in Cambridge (partially in joint work with P.L. Falb and M. Freedman, and independently by Sandberg), and even H^∞-theory, can be traced to this pioneering thesis. This tradition of research in input-output stability theory was to be continued by Jan C. Willems in his doctoral dissertation written under the direction of Roger Brockett, and culminating in his book *The Analysis of Feedback Systems* completed in 1970. Equally important was the work of Roger Brockett and his students on frequency domain stability Theory (both from a state-space and input-output points of view) with connections to the work of Popov, Kalman and Yacuvobich. These research activities led to the development of an innovative course on non-linear systems and stability theory.

During the latter part of the sixties the research excitement continued with the addition of new faculty and a new generation of graduate students. New faculty members included Fred C. Schweppe, moving from Lincoln to the campus (who strengthened the curriculum in the area of estimation and identification theory and the application of system-theoretic concepts to large-scale power systems), Jan C. Willems after he received his Ph.D. at MIT, Ian B. Rhodes after he received his Ph.D. from Stanford (who initiated research in differential games, large scale systems, and stochastic control), and Sanjoy K. Mitter from Case Institute of Technology (whose research in optimization and distributed parameter systems significantly strengthened the research carried out in the Laboratory). Unfortunately, during the same time period G. Zames and R.W. Brockett left MIT.

In addition to the development of a new curriculum in control at the graduate level, the Laboratory was involved in some significant educational efforts at the undergraduate level. As early as 1966, Professors M. Athans and M.L. Dertouzos joined Professors R.N. Spann and S.J. Mason in developing a new two-semester experimental subject at the sophomore level. These two subjects exposed the student to systems concepts using electric circuits and other dynamic systems as the main set of examples. The notion of state was exposed for both conventional electrical circuits and dynamical systems as well as for digital computer systems consisting of flip-flops and logic gates. This course sequence was quite popular, especially with students interested in computer science; over eighty undergraduates per year selected this route, and many of them were attracted to the systems and control area because of this exposure (one of the more well known students was John C. Doyle). This course sequence was terminated when the department developed the core EECS curriculum. The outcome of this educational innovation was two books co-authored by Athans, Dertouzos, Spann, and Mason, *Systems, Networks and Computation: Basic Concepts* and *Systems, Networks* and *Computation: Multivariable Methods.*

Another educational experiment that was tried out in the late sixties by Professors Gould, Rhodes, and Athans, was to teach a senior level elective in control that truly blended the modern state-space time-domain ideas with the more classical frequency-domain concepts of classical control. This educational experiment was more or less a "bust." We simply did not know enough about frequency-domain properties of multivariable systems to carry out the required intellectual unification in the eyes of the students (it took the field about ten more years to develop the necessary machinery to successfully blend these ideas).

The Period of Transition

We are now in the beginning of the seventies. The Control Group at the electronic Systems Laboratory consists of Michael Athans, Leonard Gould, Sanjoy Mitter, George Newton and Jan Willems. Fred Schweppe, although formally not associated with the Laboratory is very much part of this group. The major ideas in optimal control, state space theory of linear systems and stability theory have almost reached a definitive stage of development. Kushner, Wonham, Duncan, Mortensen, Zakai, Fleming and others, borrowing deep ideas from the theory of Stochastic Differential Equations and Markov Diffusion Processes have laid down the foundations of Non-linear Filtering and Stochastic Control. Pravin Varaiya of the University of California, Berkeley, spent the year 1974-75 in the Laboratory and gave a course on these topics. It was natural that research in the Laboratory and elsewhere should try to shift to more fertile grounds. Optimal Control theory was to move to the study of infinite-dimensional systems (delay and distributed parameter systems). Linear Systems Theory became more abstract and algebraic under the influence of R.E. Kalman and an attempt was made to unify Automata Theory and Linear Systems Theory using the theory of Rings and Modules. Roger Brockett, while still at MIT initiated a theory of non-linear control using differential-geometric ideas. These developments are reflected in the research of the control group in the Laboratory, especially in the doctoral dissertations of D.P. Bertsekas, G. Blankenship, C-Y Chong, J. Davis, A.A. Desalu, A.E. Eckberg Jr., H.P. Geering, J. Gruhl, P.K. Houpt, L.L. Horowitz, T.L. Johnson,K.M. Joseph, L.C. Kramer, A. Lopez-Toledo, D.N. Martin, T.L. Niemeyer, R.S. Pindyck, L.W. Porter, G. Prado, N.R. Sandell, Jr., M. Telson, E. Tse, H. Vandevene, M.E. Warren, D. Willner, A.S. Willsky, J.D. Wood, W. Hayes, S. Greenberg, K. Glover, and R.D. Johnston come to mind. We single out some of these dissertations to highlight the importance of their contributions. D.P. Bertsekas' work on estimation and control in the presence of unknown but bounded disturbances would in some sense be rediscovered twenty years later in the context of H^∞-estimation. G. Blankenship extended input-output stability theory to stochastic systems. A.E. Eckberg's dissertation (never published) on algebraic theory of linear systems was a forerunner of Fuhrmann's later work on the same topic. R.D. Johnston in his dissertation (also never published) on systems over rings independently proved the realization theorem for linear systems over rings and gave applications to coding theory. K. Glover undertook a fundamental study of identification in his thesis. William Hager's dissertation on Convex Optimal Control Problems used mathematical programming ideas to study state-constrained control problems. E. Tse studied optimal control problems with incomplete information. N. Sandell's dissertation was concerned with optimal control of finite-state finite-memory systems. A.S. Willsky studied non-linear filtering problems using lie-algebraic and differential geometric methods. An impressive list indeed!

It also became obvious that the theory of single-agent optimal control and estimation theory was reaching a certain state of maturity. Thus, from a theoretical perspective the efforts of students and faculty shifted to minimax problems, differential games, stochastic control with non-classical information patterns, and decentralized control theory. The famous Witsenhausen "counterexample" pointed out the clear conceptual and computational difficulties associated with the so-called nonclassical information patterns, key ingredients of any truly decentralized optimal control formulation. Theoretical research along these lines initiated at that time continues to this date.

Applications Projects in the Early Seventies

The theoretical research in large-scale systems was complemented by more applied studies. One important area during that time period was the study of safer and more automated air traffic control systems. Marc Connely was involved for several years in the evaluation of a cockpit air traffic situation display using a mockup of a Boeing 707 airliner, in collaboration with faculty from the MIT Flight Transportation Laboratory, as well as the initial efforts of the MIT Lincoln Laboratory in this area under the leadership of H. Weiss. This effort was complemented by studies of optimal flow and aircraft merging and control strategies by M. Athans, M. Telson, A.H. Sarris and L.W. Porter in the near terminal area. Some of these early (and then controversial ideas) are reaching fruition at the present time thanks primarily to extensive subsequent R&D work at the MIT Lincoln Lab and at the NASA Ames Research Center.

Another interdisciplinary project that started in the early seventies involved a collaborative effort with faculty and students in the Department of Economics and the Sloan School of Management. It started when R.S. Pindyck, an EE graduate student, decided to pursue a dual thesis involving the application of Linear-Quadratic Optimal Control methods to macroeconomic models. His Ph.D. thesis, under the supervision of Professors E.S. Kuh and M. Athans, was the first step in this collaboration. Pindyck joined the Sloan School faculty, and then together with Kuh and Athans participated in an NSF-funded effort involving the application of optimal control to economic systems. Since Professor Kuh was also the director of the Cambridge branch of the National Bureau of Economic Research (NBER), this joint research effort brought together a number of researchers from the systems and economics discipline. In addition to Kuh, Pindyck and Athans, the following participated in several ways in the research: T.L. Johnson, D.Kendrick, K. Wall. A.H. Sarris, T. Ozkan, M. Szeto, J.W. Neese, M. Telson, E. Tse, V. Klema, L. Papademos. A large scale econometric model of the U.S. economy, involving hundreds of state variables, and several control variables was used to understand coordinated monetary and fiscal control issues. A series of workshops, co-chaired by M. Athans and G. Chow (from Princeton), were organized and attended by control theorists and economists alike, resulting in cross fertilization of ideas especially in the areas of system identification and stochastic adaptive (dual) control.

Chapter 3

From the Seventies to the Present: From Electronic Systems to Information and Decision Systems

Transition and Change

In 1973 Professor Michael Athans was appointed Director of the Electronic Systems Laboratory and the Laboratory was formally designated an interdisciplinary Laboratory reporting to Provost Walter Rosenblith. During the same time frame several other changes were taking place. With the appointment of Professor W.B. Davenport as head of the EECS department, the split between EE and CS was avoided. Area I and VI joined into a new EECS graduate area in Systems, Communications and Control with Sanjoy Mitter serving as the first chairman of the new area. With the encouragement of Professor Davenport, Professors J.M. Wozencraft, R.G. Gallager and A.W. Drake joined the Laboratory. We also added several new young faculty during these years: T.L. Johnson, A.S. Willsky, N.R. Sandell Jr., C. Leung, A. Segall, and later on in the late seventies, D.P. Bertsekas, P. Humblet, B. C. Levy, G. Verghese, and R.R. Tenney. Alan Willsky was appointed Assistant director of the Laboratory in 1974 and Robert Gallager, Associate Director in 1975. The Laboratory acquired its present name in 1978 to reflect the broadening of its horizons.

From an academic point of view the emphasis in control theory continued in the Laboratory, slanted to the areas of distributed parameter systems, decentralized control, differential games, algebraic systems, and control system design. Michael Athans produced 70 TV video tapes on Modern Control Theory, together with associated study guides and CAD software (written by N.R. Sandell, Jr.), which were distributed by the MIT Center for Advanced Engineering Study. These tapes were used extensively by industry and universities in the U.S. and abroad. It is only during the last year or so that the demand for these tapes has diminished. New subjects reflecting new research directions in communication networks and stochastic systems were introduced. From an intellectual point of view we were witnessing a unification of concepts and theories in systems, stochastics, communications, and operations research. This was reflected in the development of a new graduate core curriculum in Systems, Communications and Control.

We continued to attract excellent graduate students whose theses pushed the state of the art. We recall the numerous research contributions in the Systems and Control area by the following doctoral students in ESL (and later in LIDS) in the mid-and late seventies: A. Akant, Y. Baram, S.M. Barta, J.D. Birdwell, D.A. Castanon, T.E. Djaferis, J.R. Dowdle, S.G. Finn, C.S. Green, P.K. Houpt, P. Kam, H. Kan, W. Kohn, R.T. Ku, J. Liu, D.P. Looze, H. Chizeck, D. Teneketzis, L.A. Monauni, P. Moroney, L. Platzman, M. Safonov, H. Sira-Ramirez, R.R. Tenney, and K. Yared, S.J. Marcus, R. Kwong, J. Wall, R.D. Washburn, J. Eterno, F. Moss, E. Chow, M. Bowles, D. Ocone, L. Horowitz, J.I. Galdos, S. Young

Communication Research in the Laboratory

The Early Years in the Research Laboratory of Electronics

Communication research has had a long and and illustrious past at MIT. Before moving to LIDS in 1975, the focus of communication research was in RLE, the Research Laboratory of Electronics. Much of the early growth of information theory took place there in the 50's and 60's. Claude Shannon moved to RLE from Bell Laboratories in 1956 and added his genius to a research group already containing Profs. Robert Fano, Peter Elias, David Huffman, and their students. Jack Wozencraft soon finished his thesis on sequential decoding and joined the faculty. In closely related work under the intellectual leadership of Norbert Wiener, Jerome Wiesner, Wilbur Davenport, Y.W. Lee, and their students were doing seminal work on detection and estimation. These research efforts were part of a larger intellectual ferment looking at the theoretical foundations of communication, computer science, artificial intelligence, and biological communication and computation.

This early era at RLE was characterized by a great enthusiasm for the value of conceptual and mathematical understanding of fundamental issues in communication, control, and computation. This enthusiasm was shared not only by the researchers themselves, but also by the MIT administration and by the funding agencies. The

Alvin W. Drake

John M. Wozencraft

George C. Verghese

intellectual excitement combined with the sense of value attracted an extraordinary group of graduate students into communication research, including Jack Wozencraft, Amar Bose, Irwin Jacobs, Barney Reiffen, Robert Gallager, James Massey, Jacob Ziv, Thomas Kailath, George Zames, Elwyn Berlekamp, David Forney, Leonard Kleinrock, Robert Kennedy, Harry van Trees, Donald Snyder, Fred Jelinek, David Sakrison, and Arthur Baggeroer.

A number of textbooks came out of MIT during this period that revolutionized the teaching of communication theory throughout the world. *Random Signals and Noise* by Davenport and Root (McGraw Hill, 1958), *Transmission of Information* by Fano, (Wiley, 1961) *Principles of Communication Engineering* by Wozencraft and Jacobs (Wiley, 1965), *Information Theory and Reliable Communication* by Gallager (Wiley, 1968), and the four volume set *Detection, Estimation, and Modulation Theory* by Van Trees (Wiley, 1968-1971) are books that are still in use today, both as texts and references.

Toward the end of the 60's, the academic climate had changed somewhat. There was a disillusionment with military support for research, and this quickly became subverted into the current idea that research should be "aimed" at immediate social and or industrial problems rather than toward understanding basic questions. There was also a widespread feeling that communication research had far outstripped applications. For these reasons, communication research in RLE became somewhat less active and somewhat more application oriented.

As interest in traditional point to point communication waned, a critical need was developing for basic research on data networks. At the same time, it was becoming increasingly apparent that the distributed nature of many large control problems required totally new approaches. At one level, there was the recognition that data networks required an integration of both communication and control ideas and that distributed control required a similar integration. At a deeper level, there was a set of research areas involving distributed algorithms, distributed estimation, and distributed computation that were recognized to be fundamental in their own right.

The Move to The Electronic Systems Laboratory

In 1974, the decision was made to start a major effort on data networks in the Electronic Systems Laboratory (now LIDS). Prof. Wozencraft returned to MIT from the Naval Postgraduate School at Monterey to take part in this effort, and Prof. Gallager moved to ESL from RLE. In addition, Prof. Adrian Segall received a faculty appointment and moved to ESL from Systems Control Inc. Profs. Athans and Sandell were also involved in the project, particularly with respect to relationships with distributed control and with military C^3 problems.

The personnel involved with communication at LIDS has changed considerably over the years. In 1977, Jack Wozencraft returned to Monterey and Adrian Segall moved to The Technion. Prof. Pierre Humblet, who had received a Ph.D. at MIT in 1977, joined the faculty in 1978 and Prof. James Massey spent a year's sabbatical at the Lab at the same time. Prof. Dimitri Bertsekas came to LIDS from the University of Illinois in 1979, and, while focussing on data networks, strengthened other areas at LIDS. Finally Prof. Kennedy, who had been working on optical communication at RLE, moved to LIDS in 1985 to start a project on optical fiber networks.

In the early days of the network project, the emphasis was on dynamic and quasi-static routing policies. This was a much more challenging problem than routing in voice networks because of the bursty nature of the traffic. The problem involved distributed control in an essential way because the delays involved in communicating control information were comparable to the traffic fluctuation times. Early efforts at optimal dynamic solutions included theses by Moss, Defenderfer, Wunderlich, and Ros Peran. Later theses by O'Leary, Tsai, and Castineyra were less concerned with optimal dynamic solutions but involved more general models.

It soon became clear to the data network group that, although routing was a fascinating theoretical problem, the primary practical problem was that of congestion control. The issue here is not only to prevent congestion in the network, but also to provide fair access to all the users even under heavy load conditions. The nature of the congestion

Alan S. Willsky, T. Michael Chin

From Left: S. Muller, Amiel Feinstein, Peter Elias, David Huffman, Saburo Muroga, Robert M. Fano

problem changes both with the mix of traffic in a network and with the higher layer protocols employed, and thus congestion control continues to be an area of research today. Early theses on congestion by Golestaani, Ibi, and Gafni focused on combined congestion and routing. Later theses by Friedman, Regnier, and Tiedemann were concerned with controlling the queues at the network nodes, while Mosely, Mukherji, and Hahne focussed on algorithms to provide fairness between the sessions.

A byproduct of the early work on routing and congestion was the realization that the design of distributed algorithms is an important research problem in its own right. The communication required by the early algorithms used in network control often far outweighed the data itself, and the need for avoiding excessive control information has persisted to the present day. The early work at LIDS on this problem has stimulated much work elsewhere on the communication complexity of distributed algorithms. The problem of designing distributed algorithms that are robust in the presence of errors and failures has been another major research topic. Roskind and Spinelli both solved important problems in the area of failure recovery.

Multiaccess communication was also a major research topic of the network group. The Capetanakis random access algorithm was developed in a 1977 Ph.D. thesis, and Gallager's first-come first-serve splitting algorithm and the improvement by Mosely and Humblet were developed soon after. Theses by Hluchyi, Helman, and W. Lee explored random access for a limited number of users and various network models. The multiaccess work has now broadened into a general investigation of the use of coding, spread spectrum, and random access for mobile and personal radio communication. Ozarow, Arikan, Hui, and Hughes have made major contributions here.

The study of data networks at LIDS had reached a fairly mature stage by the mid 80's, and the textbook *Data Networks* by Bertsekas and Gallager (Prentice Hall 1987) was intended to provide a synthesis of understanding to a field where knowledge had been rather fragmented.

Estimation Theory, Signal Analysis and Inverse Problems

Estimation theory and the related fields of signal analysis and inverse problems have been topics of research in the Laboratory since the late sixties but have been vigorously pursued since about 1973. This continues the tradition of Wiener's fundamental idea of extracting probabilistic information from signals. This research has been pursued by Sanjoy Mitter, Alan Willsky, and later by Bernard Levy, John Tsitsiklis, George Verghese and their graduate students. Alan Willsky wrote his thesis on non-linear filtering and considered problem formulations in which the state space was a manifold rather than just Euclidean space. These problems were motivated by problems in analog communication such as phase de-modulation. This work was later continued by S.J. Marcus in his doctoral dissertation. Willsky showed (partially in joint work with J.T. Lo) that certain non-linear filtering problems have finite-dimensional filters. In the seventies he also initiated a rapprochement between control and filtering theory based on state-space methods and problems of signal processing. This research led to his book *Digital Signal Processing and Control and Estimation Theory: Points of Tangency, Areas of Interaction and Parallel Directions*, MIT Press, 1979.

The rapprochement with signal processing developed through intensive interactions with Professor A. Oppenheim. The fact that systems and control ideas have an important role to play in Signal Processing dates back to this period. A by product of this interaction was the collaboration of Oppenheim and Willsky in teaching *6.015* and the writing of the widely-used undergraduate text *Signals and Systems.*

Ultimately this work was to develop into the notion of model-based signal analysis, an area in which the Laboratory has continued to be active. A related activity of Alan Willsky's on failure detection using generalized likelihood ratio test ideas. Sanjoy Mitter also became interested in non-linear filtering and first pursued ideas for obtaining lower bounds for estimation errors using ideas of rate distortion theory. Subsequently in joint work with D.

Allinger he was able to solve one of the major open theoretical problems in non-linear filtering - the so-called innovations conjecture due to Frost and Kailath. He and, independently, Roger Brockett suggested a new attack on the non-linear filtering problem using the Duncan-Mortensen-Zakai equation and Lie algebraic ideas especially with a view to obtaining finite-dimensional non-linear filters. These ideas led to a flurry of activity in many places even though most of the results turned out to be negative. Doctoral students who wrote important theses on these topics include L. Horowitz, J.I. Galdos, R.H. Kwong, S.I. Marcus, P. Kam, L. Platzman, J.E. Wall, S. Young and D. Ocone.

As the decade of the seventies was coming to an end it had become clear that progress in non-linear filtering and stochastic control was likely to be difficult without some major new idea. The group shifted its attention to multi-dimensional signal analysis, motivated by problems of image analysis and vision, tomography, geo-physical structure determination, and most recently x-ray Crystallography. The major difference between signals which depend on time and signals which depend on space or space-time is that there is no natural ordering for spatial signals. Modeling of one-dimensional signals via stochastic processes has to be replaced by the far more difficult idea of modelling via a random field. The incorporation of a priori knowledge is also a more complicated matter and geometry enters in an essential way into the modelling process. An early thesis on random fields was that of R. Washburn.

The new conceptual problem of concern is the determination of structure from noisy and sparse measurements of spatial (or spatio-temporal) signals. The contributions of this group have been manifold. We only cite a few to give the flavor of these contributions. Levy and his students, using ideas from inverse scattering, signal processing and tomography, solved a variety of inversion problems - for example, joint reconstruction of velocity and density in an elastic medium. Willsky and his students have investigated a number of geometrical reconstruction problems - for example, determining a convex set from noisy support plane measurements and devising optimal methods for detecting, localizing and estimating the shape and orientation of objects from tomographic measurement. Mitter and his students investigated problems in computer vision, such as depth from stereo and boundary detection using ideas of Markov random fields on a lattice. The best measure of this contribution are the doctoral dissertations written by students of this group. We cite the dissertations of M. Adams (Estimation for Non-Causal System Models), M. Bello (Map Up dating and Terrain Analysis), E. Chow (Failure Detection) M. Coderch, X. C. Lou, R. Rohlicek (Multiple Time Scale Analysis of Linear Systems and Markov and Semi-Markov Chains), D.J. Rossi (On Tomographic Reconstruction), Y. Avniel (Scattering View of Stochastic Processes), P.C. Doerschuk (Markov Chain Approach to Electrocardiogram Analysis) J.P. Greshak, R. Nikoukah (System and Estimation Theory for Non-causal Systems) J.L. Prince (Geometric Reconstruction of Convex Sets and Tomographic Reconstruction from Sparse Data), J.L. Marroquin (Probabilistic Solution of Inverse Problems), M.B. Propp (Thermodynamic Properties of Markov Processes), A.E. Yagle, C. Esmersoy, A. Ozbek, C.D. Banks (Inverse Problems Arising in Geophysics), R.H. Lamb (Parametric Non-linear Filtering) A.H. Tewfik (Spectral Estimation for Random Fields) and G.W. Hart (Minimum Information Estimation of Structure).

A new approach to x-ray crystallography which uses Markov random fields to incorporate a prior chemical knowledge and bypasses phase reconstruction, has been developed by Peter Doerschuk, Sanjoy Mitter and Alan Willsky. Mitter, his students, Richardson and Kulkarni and Ofer Zeitouni (Technion and LIDS) have investigated problems in computer vision such as boundary detection using variational calculus, geometric measure theory, and Markov random fields. Work on motion estimation has been done by Willsky, Levy and their collaborators.

The final major topic in this circle of ideas, which brings us to the present, is that of multi-scale stochastic processes. Motivated by developments in the theory of wavelets, currently a hot topic, Alan Willsky and his students, in collaboration with Albert Benveniste of INRIA, France, have undertaken a systematic study of multi-resolution description of stochastic processes described as evolving on trees and lattices and related problems of estimation. It would appear that a reasonably complete estimation theory, with its algorithmic counterpart such as Schur and Levinson recursions, can be constructed for this class of problems. This research should be important for the study of stochastic fractals and multi-scale signal inversion problems.

The Interaction of Theory and Applications

The period from the mid-to late seventies was one of rapid changes in more applied projects in the Laboratory. With Project INTREX winding down, we saw a large turnover in the research staff and programming staff employed by the Laboratory. New projects, many based on theoretical concepts only a few years old, were started and new research staff members were appointed together with the Laboratory faculty and students. During that time period, the following were hired as Laboratory staff members: Dr. K-P Dunn, Dr. D.A. Castanon, Dr. P.K. Houpt, Dr. S.B. Gershwin, Ms. I. Segall, Dr. B.C. Levy, Dr. A. Laub, Ms. V. Klemma, Dr. A.H. Levis, Ms. E.R. Ducot, and Dr. L. Valavani.

One of the first applied projects was related to adaptive control and failure detection algorithms for the NASA F-8 aircraft done for the NASA Langley Research Center. NASA was using the F-8 as a digital fly-by-wire (DFBW) testbed and they wanted to investigate the applicability and performance of advanced adaptive flight control and failure detection algorithms. The adaptive control method we investigated was the so-called Multiple Model Adaptive Control (MMAC) rooted in dynamic hypothesis testing problems. This algorithm was examined by several researchers and in particular in the doctoral thesis of Willner in 1973 under the supervision of M. Athans. The F-8 MMAC project involved several faculty, staff, and students (M. Athans, N.R. Sandell, A.S. Willsky, K-P Dunn, D.A. Castanon, I. Segall, C.S. Greene, W-H Lee, Y. Baram) and involved specific digital adaptive designs tested at numerous simulations (with some hardware in the loop) at NASA Langley. It turned out that the MMAC algorithm was too unpredictable, and NASA wisely decided not to flight test our MMAC algorithm at Edwards AFB. However, we learned a lot about shortcomings of optimal control designs, and LQG compensators in particular, and learned to be suspicious of practical adaptive algorithms. This research project motivated two separate subsequent Ph.D. theses (by Y. Baram and C.S.Greene) to understand issues in MMAC from a deeper perspective. Another side benefit of the F-8 project was that we established a close relationship with Dr. Gunter Stein of Honeywell (who headed a parallel Honeywell effort that did get flight tested). Stein was appointed as Adjunct Professor in the Electrical Engineering and Computer Science Department and continues to be a key member of the LIDS Control Group.

The area of research in failure detection represents a good example of the interactions between theory and practice. It arose out of Willsky's work on inertial system calibration while he was a consultant at TASC. This led to his research on the theoretical aspects of failure detection which in turn led to i) the successful flight-testing of the dual redundancy failure detection system for the F-8 aircraft and ii) a system for detecting arrythmias in electrocardiograms. Chizeck's thesis on fault-tolerant optimal control is an outgrowth of these ideas.

Another applied control and estimation project involved the automated detection of traffic incidents on freeways and dynamic control of freeway traffic via ramp metering, especially in freeway-corridor systems (i.e., situations where several dense origin-destination nodes are interconnected by at least two major freeways or large arteries). This research project was funded by the U.S. Department of Transportation and the Federal Highway Administration, and won on the basis of open competitive bidding. It attracted a large group of faculty, staff and students (P.K. Houpt, S.B. Gershwin, M. Athans, A.S. Willsky, J. Ward, E.Y. Chow, C.S. Greene, A.L. Kurkjian, D.P. Looze, D. G. Ohrlac, H.N. Tan, W. Mitchell, J. Olesik, R. Lopez-Lopez). From a technical point of view, generalized likelihood ratio techniques and multiple-model hypothesis testing methods were used for incident detection. Large-scale mathematical programming concepts were used to optimize the steady-state traffic flow, and variants of LQ and LQG theories were used for ramp-metering control strategies.

Starting in the mid-seventies and continuing to the early eighties the U.S. Department of Energy and the Office of Naval Research sponsored several research projects in the laboratory in the areas of large scale systems and of decentralized estimation and control. These projects allowed Alan Laub and Virginia Klema to join the research staff of the Laboratory and establish research activity in numerical methods for estimation and control. One key important outcome of this research was the development by A. Laub of a new algorithm for solving high-order algebraic Riccati equations. Laub's algorithm is at the heart of present day commercial software packages (e.g. CTRL-C, MATRIX-X, etc) in the control and estimation areas. From a theoretical point of view, we learned many more intricacies related to the control of large scale systems, as manifested in the Ph.D. theses of D.A. Castanon, S.A. Barta, W. Kohn, D.P. Looze, J.E. Wall Jr., R.R. Tenney, and D. Teneketzis. We were fortunate to have Professor P.P. Varaiya from the Univ. of California at Berkeley and David G. Luenberger from Stanford spend their sabbatical leaves at LIDS.

Significant advances were also being made in the field of multivariable control. To a large extent, some of the key questions arose from problems associated with Linear-Quadratic (LQ) and Linear-Quadratic-Gaussian (LQG) designs in the presence of multiple saturations and significant uncertainty (both the F-8 project and the freeway corridor project contributed pragmatic motivation). Guaranteed gain-margin robustness properties for multivariable LQ designs were first derived by P.K. Wong in his SM thesis (1975) using geometric arguments based on Wonham's work; his work was continued by Safonov who presented a unified approach to stability-robustness and multivariable gain and phase margin properties of LQ and LQG designs in his Ph.D. thesis (1977), under the supervision of M. Athans, published later as an MIT Press research monograph. Safonov's seminal thesis work opened up brand new directions in frequency-domain oriented stability-robustness results and spearheaded subsequent research in that particular area both at MIT and elsewhere, notably the research of J.C. Doyle. The role of Gunter Stein during that time period as a catalyst between the MIT and Honeywell groups (and others) was indeed critical, and led to a series of key papers co-authored by Doyle and Stein on the role of the singular value decomposition in frequency-domain synthesis of multivariable control loops, the LQG/LTR method, etc. Other contributions in that time period in the design of control compensators can be found in the doctoral theses of J.D. Birdwell, T.E. Djaferis, J.R. Dowdle, R.T. Ku, and P. Moroney. The work of Moroney represents an application of digital signal processing ideas to the digital implementation of control compensators.


Gunter Stein


Wilbur B. Davenport


Alexander H. Levis

Another project that was started in the mid-seventies dealt with system-oriented aspects of automated manufacturing. The research was initially funded by the NSF Research Applied to National Needs (RANN) initiative and we were able to negotiate a large scale multiyear award in this area. A large number of researchers contributed to this effort (S.B. Gershwin, J.E. Ward, M. Athans, D.P. Bertsekas, J.N. Tsitsiklis, P. Kanallekis, E.L. Hahne, J.G. Kimemia, M.H. Ammar, and others). This project led to significant interactions with industry. It was hoped that flexible manufacturing systems could be viewed dynamically, and that the hierarchical planning of production in a factory was amenable to quantitative studies. The project served as a melting pot for integrating static and dynamic optimization, queueing theory, scheduling and stochastic control. The complexity (NP-completeness) of many important problems in manufacturing became evident as the research progressed. This research would continue well into the eighties with support provided by the U.S. Army and IBM, and would lead to the Hierarchical Scheduling Paradigm for Flexible Manufacturing Systems. Many investigators, such as P. Kumar (University of Illinois) would be inspired by this research.

Soon after the communications research group was started in LIDS, Professors Athans, Davenport, Wozencraft and Sandell decided to explore large-scale system and communication issues in the general area of military Command, Control, and Communications (C^3) systems. From a pure theoretical perspective, prior research in distributed decision-making in large scale systems always involved awkward tradeoffs regarding the amount of information that had to be shared. The complexities of nonclassical information patterns and signalling phenomena were appreciated since the late sixties. Military C^3 systems were appealing as research paradigms not only because of their inherently dispersed nature, but also because the consequences of too much communication were clear: the enemy could intercept radio communications and could then identify and locate. With generous support from the Office of Naval Research and the Air Force Office of Scientific Research, starting in 1977 the research of the LIDS/C^3 group was initiated. This research continued until 1990. Each summer we sponsored the annual MIT/ONR C^3 Workshop, which attracted a dedicated and diverse group of researchers in the field. This workshop continued for nine years, and was then replaced by the on-going Joint Directors of Laboratories (JDL) annual C^3 symposium. As the project matured more and more researchers contributed to the MIT research such as A.H. Levis, E.R. Ducot, D.A. Castanon, R.R. Tenney, J.T. Casey, and A. Segall as well as several students (K.M. Keverian, R.H. Lamb, Jr., S. Andreadakis, K.L. Boetcher, C. Bohner, D. Perdu, R.P. Wiley, M. Ma, P.J.F. Martin, S.A. Hall, H.P. Helion, J. Kyratzoglou, A. Louvet, F. Valraud, F.E. Bruneau, R.C. Magonet-Neray, C. Mok, V.O.K. Li, K.T. Huang, L. Ekchian, P.A. Remy, W.C. Roth, A. Ozbek, G. Chyen, P.H. Cothier, E.A. Hinzelman, P.P. Ng, C. Lee, J.D. Papastavrou, J.L. Grevet, P.A. Hossein, J. Walton, G. Polychronopoulos, S.T. Weingaertner, V.Y.Y. Jin, and others).

Significant new system-theoretic problems arose from this work. In the surveillance area, under the leadership of Sandell and Tenney, the distributed detection problem became a standard paradigm for truly distributed decision-making with communication constraints. This original idea was investigated in four subsequent doctoral theses at MIT and opened up a new field of research. Also in the surveillance area, the problems associated with data association and the multi-target multi-sensor tracking problem received a great degree of attention and novel algorithms (now being routinely used in advanced defense systems) were developed. Davenport and his students examined the blending of distributed data bases within radio communication networks. Gallager, Segall, Wozencraft, and Bertsekas led the effort in the development of fail-safe routing algorithms in vulnerable communication networks. Finally, Levis, Tenney, Tsitsiklis, and Athans approached the organizational design problem using a variety of normative and descriptive approaches (information-theoretic models, Petri net

Lena Valavani

Munther A. Dahleh

Pierre A. Humblet, Robert S. Kennedy

technology, distributed optimization algorithms, etc.). In short, the C^3 LIDS project contributed a large variety of concepts, ideas, and formulations to the field; in the process it opened up brand new research directions in the area of large scale systems, and organizational design.

Systems and Control Research: 1980-1985

The newly developed ideas in blending time-domain and frequency-domain concepts (such as singular value loop shaping) in multivariable control synthesis which originated from the work of Safonov and Athans, and Doyle and Stein in the late seventies continued strongly in the 1980-1985 era. Several doctoral students made significant contributions to the field: S.M. Chan, P. Dersin, N. A. Lehtomaki, J.N. Tsitsiklis, C.E. Rohrs, M. Shahjahan, P.M. Thompson, and D.M. Orlicki. During that period of time the newly developed methods were evaluated in a series of applications-oriented problems such as the control of multi-terminal DC networks (jointly with GE) which involved the decentralized control of a dynamic model of power generation for the Western US (Athans, Levy, Petkovski, Chan, Lehtomaki, Ng), automotive engine control (Lee, Lewis), the multivariable control of jet engines using models of the F-100, GE-21 and TF-700 (Kappos, Idelchick, Kapasouris, Pfeil, Dunn), several aircraft and helicopter control oriented problems (MacMuldroch, Bodson, Haiges, Quinn, Rodriguez), submarine control (Lively, Mette, Martin) to mention just a few. G. Stein continued his affiliation with LIDS as an Adjunct Professor. The control group was strengthened by the appointment of A.H. Spang, III, from GE, also as an adjunct professor, who provided key leadership in the area of jet engine control. M. Athans and G. Stein collaborated in revising the graduate level EECS subject (6.232) to reflect the new way of thinking and the new results in control system design. In the adaptive control area a number of students (Rohrs, Krause, Orlicki, LaMaire) under the supervision of Valavani, Athans and Stein demonstrated the pitfalls of adaptive control in the presence of high-frequency unmodelled dynamics. C.H. Rohrs' Ph.D. thesis opened up a new area of research on robust adaptive control, which is still receiving a great deal of attention by the adaptive control community. Professor S.S. Sastry joined the faculty in this period and strengthened our research in non-linear and adaptive control.

Control: 1985 to the Present

During the past few years the educational and research innovations in the control area have continued at a fast pace. Professor M.A. Dahleh joined the EECS faculty and LIDS. Once more we have had excellent doctoral students during this period, some still in the pipeline (I. Diaz-Bobillo, D. Flamm, D.G. Grunberg, P. Kapasouris, R. LaMaire, W-H Lee, D. Milich, D. Obradovic, B. Ridgley, A.A. Rodriguez, J.S. Shamma, P. Voulgaris, and others).

From an educational perspective the EECS, Aeronautics and Astronautics, Mechanical Engineering, Chemical Engineering, and Ocean Engineering Departments, agreed to offer a joint interdisciplinary two-semester graduate subject on multivariable control systems. The following faculty from these five departments participated in the development and teaching of this popular subject: Professors M. Athans, M.A. Dahleh, K. Hedrick, G. Stein, G. Stephanopoulos, M. Triantafyllou, L. Valavani, W. Vander Velde, and B. Walker. This subject sequence contains a heavy dose of computer-aided design problem sets and relies upon the resources of Project ATHENA. Typically 70-90 graduate students start taking the subject, and about 35-55 finish it in the spring. The interaction of the faculty and graduate students from these different departments has had a most positive impact in the research at both the SM and Ph.D. levels. New courses on linear and non-linear systems which reflect ideas of uncertainty reduction and robustness have been introduced.

The raison d'etre for feedback control is to combat uncertainty. This was well understood in the fifties but somewhat forgotten in the sixties. It reemerged in research in the Laboratory in the seventies, but its decisive resurrection can be attributed to the seminal paper of

Zames in 1981 where he formulated the H^{∞}-problem of sensitivity minimization and later solved it in joint work with Francis and Helton.

Sanjoy Mitter and his student David Flamm provided one of the earliest solutions to the problem of minimizing the H^{∞} norm for infinite-dimensional plants (e.g. plants with delays). Under the supervision of Dahleh, Stein and Mitter, LIDS contributed in many different ways to the advancement of the H^{∞}-theory. Recently, Armando Rodriguez, a student of Munther Dahleh provided a methodology for designing finite-dimensional controllers for infinite-dimensional plants via approximation. On the other hand, Dahleh and his students have been involved in developing the l^1 design design methodology, a topic he pioneered, which is a worst case design methodology in the presence of bounded, persistent but unknown disturbances. It attempts to achieve time-domain specifications in the presence of plant uncertainty.

In the past five years most of the research in the control area has revolved in the area of robust control for both linear and non-linear systems. In the linear systems area Dahleh and his collaborators have done research to understand the robustness of l^1 methodology.In the non-linear systems area significant progress has been made in extending the loop transfer recovery concept's to non-linear designs (Grunberg), and in dealing with multivariable magnitude and rate saturations (Kapasouris), and theoretical issues associated with gain scheduling (Shamma).

With the growing need to understand the fundamental limitations and capabilities of controller design in the presence of uncertainty, more attention has been given to the problem of non-parametric identification. A joint effort by Dahleh and Tsitsiklis on the study of this problem has been initiated at LIDS. The objective is to derive a theory for systems with uncertainties that makes it possible to design robust controllers with learning features.

Finally, a focus of a growing size in the control community is the analysis of Discrete-Event Dynamic Systems (DEDS), i.e. complex, usually man-made systems, whose significant dynamic behavior can be described by the sequential occurrence of discrete events. Flexible manufacturing, interconnected power networks, and distributed computer systems are but three examples. A recent focus of much of this research, introduced by Prof. W.M. Wonham of the University of Toronto, attracting the interest of numerous researchers including Varaiya at Berkeley, is in the development of a qualitative theory of control for such systems using constraints and techniques opted from such computer science domains as automata, language theory, and temporal logic. In their work in this area Willsky and his student Cüneyt Ozveren focused on developing the elements needed to establish a true servo theory for such systems. Central to this is a notion of stability which, more accurately, corresponds to a property of error-recovery or resiliency more closely allied to the notion of catastrophic error propagation in convolutional coding. Their work also highlights the important issue of the timing of information and control. However, as Tsitsiklis has pointed out, many problems in this area suffer the curse of NP-completeness and this provides more than ample motivation for the continuing effort aimed at the incorporation of structure and more powerful modeling techniques from computer science in order to develop the techniques of control so greatly needed in this age of proliferating automation.

Communications Research: 1985-Present

In the mid 80's, rapid progress in optical fiber technology was causing a revolution in communication network research. Before this time, link capacity was a scarce resource in networks and thus efficient link utilization was a primary objective of network research. In addition, typical bit error rates on communication links were on the order of 10^{-6}, making reliable network communication a non-trivial task. With single mode optical fiber links, however, link capacities of many terabits (10^{12} bits) per second are possible in principle and bit error rates are almost negligible. Unfortunately, it appears that most processing at the network nodes will have to be done electronically well into the future, leading to the so-called bottleneck at the electronic- optical interface.

In 1985, a major new project to study such problems in the context of local area networks was started. Prof. Kennedy, who had been doing research on optical communication in RLE, moved to LIDS and Prof. Humblet, who had been working on wide area networks at more conventional data rates, soon joined the project. The original focus of their research was to investigate the potential physical layer and multi-access structure of local area networks in the next 10 to 20 years, when link capacities on the order of terabits per second are expected to be available. The challenging questions here are how to utilize the multiaccess capabilities of fiber to allow simple electronic processing at the sources and destinations. Initial investigations concerned wavelength division multiplexing, tunable lasers and detectors, optical couplers, etc. Early Ph.D. theses in this area were done by Wagner, Abernathy, Liew, Escobar, Wong, and Wasem.

Currently, work is continuing on the physical layer aspects of local area networks, but there is also increasing interest both in wide area optical networks and in the higher layer issues of congestion control, internetworking, and overall system architectures. There is an increasing recognition that these technical issues can not be separated from the questions of potential applications of optical networks, from regulatory issues, and from the economic issues of who builds such networks and how the services are tariffed. Studying these issues requires a broad interdisciplinary effort, including large scale experimentation and including extensive interactions with industry. There is a serious effort currently to integrate the optical network work in LIDS with Lincoln Laboratories (which is capable of experimental work on a larger scale than appropriate at LIDS), with other laboratories at MIT, and with industry. The current work on multiaccess radio networks will fit into this larger effort as part of the internetworking problem.

Research on Computation

Background

Many computational algorithms arising in the context of communication and control have been investigated in the Laboratory over the years. It became clear many years ago, however, that together with specific application-oriented algorithms, it was important to focus some research on broad methodological issues of computation, particularly as it related to optimization.

Optimal control problems provided a strong computational challenge in the sixties. This important class of large, computationally intensive problems taxed the memory and processing capability of the existing digital computers, and became a vehicle for much research on descent optimization methods, two-point boundary value problems, dynamic programming, and other techniques. Athans devoted substantial attention to computational methods in his optimal control course, and several theses from the late sixties contributed to the state of the art in numerical optimal control. Mitter, whose background included important contributions in the development of Newton's method and the conjugate gradient method for continuous-time optimal control, joined the Laboratory in 1970 and strengthened substantially the effort in this area.

Parallel to, and independently of, the computational research in optimal control, there were important developments in mathematical programming within the then growing operations research community. It became apparent in the early seventies that there were strong connections between mathematical programming and optimal control; at a certain level of mathematical abstraction they were equivalent, and they could be viewed within a common framework. This realization was reflected in the curriculum when Shapiro and Magnanti (from the Sloan School) and Mitter established two joint courses between Electrical Engineering and the Sloan School, focusing on linear and non-linear programming, optimal control, and the associated role of convex analysis. The convergence of the optimal control and mathematical programming methodologies was also reflected in research at the Laboratory. The theses of W. Hager and S.I. Young reflect this development. In the early sixties Bertsekas and Mitter developed the e-subgradient method that was motivated by Luenberger's research on optimal control problems with kinks (at Stanford), but also turned out to be the starting point for much subsequent work on nondifferentiable non-linear programming.

Later Years

The research effort on computational optimization was strengthened in 1979 when Bertsekas returned to the Laboratory after spending eight years at Stanford University and the University of Illinois, working mostly on non-linear and dynamic programming. A course was then established on dynamic programming and stochastic optimal control, which emphasized traditional control theory subjects, as well as a broad variety of applications in engineering and operations research. Some of these application areas, such as manufacturing systems and communication systems, were being actively researched in the Laboratory. The 1987 book on *Dynamic Programming* by Bertsekas reflects some of the related activity. At the same time, there was much research on constrained optimization, and particularly Lagrange multiplier methods and gradient projection methods (both important methods for large-scale optimal control and data network routing problems). This research was the subject of the 1982 research monograph by Bertsekas and several papers by Bertsekas and Gafni.

In the late seventies, a substantial program in large-scale network optimization was also initiated, motivated by optimal routing problems in data networks already investigated by Gallager, and by large-scale linear and non-linear optimization problems arising in communication, transportation, and manufacturing contexts. The research on network optimization continued through the eighties and resulted in new methods (based on the notions of auction and e-complementary slackness) that are radically different from the traditional simplex and primal-dual methods. As a byproduct of this research, several state-of -the-art network optimization software packages (RELAX, AUCTION, MULTIFLO) have been written and have been distributed to hundreds of users in academia, government, and industry. The Ph.D. research of Gafni, Luque, Tseng, and Eckstein have contributed substantially in this effort. Tseng returned to the Laboratory in 1987 and in partial collaboration with Bertsekas, Luo, and Tsitsiklis, made many contributions to a variety of large-scale optimization methods, including interior point methods for linear programming.

Two new and important computational research trends emerged in the late seventies and intensified during the eighties: the analysis of distributed systems, and the application of computer science (complexity) methods in communication and control.

The intensive research program on data communication networks, naturally led Gallager, Humblet and Segall to the study of distributed algorithms for routing and several related problems that arose in this context. These algorithms had to be implemented on-line and practical considerations led to the study of asynchronous distributed algorithms. Research on asynchronous algorithms, carried out initially by Bertsekas and later by Tsitsiklis, eventually

John N. Tsitsiklis

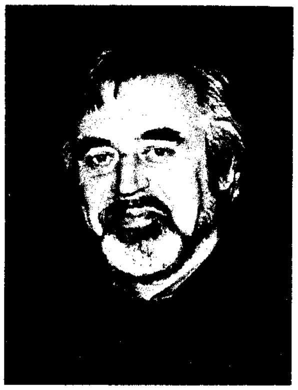

Dimitri P. Bertsekas

expanded in scope and went beyond the data network context. For example, it was found to be relevant to distributed decision making, within the context of research on C^3 systems headed by Athans. By the late eighties, the asynchronous distributed versions of most major classes of numerical algorithms relevant to optimization and control had been thoroughly investigated, resulting in a comprehensive convergence theory for asynchronous computation. A major part of the 1989 book *Parallel and Distributed Computation: Numerical Methods* by Bertsekas and Tsitsiklis was an exposition of this theory. The theses of Mosely and Tsai were also in this area.

During the eighties, it also became clear that large intelligent systems would have very intensive computational requirements, and that only massive parallelization could offer a hope of meeting these requirements. Research was carried out by Bertsekas, Castanon, Eckstein and Tseng on parallel algorithms for network flow problems, linear and non-linear programming (decomposition methods), and by Levy and Kuo on parallel multigrid algorithms for partial differential equations.

In addition, the investigation of spatial and spatio-temporal estimation and signal processing problems raises new and substantially greater computational problems than their time series counterparts. Indeed the notion of recursion in such processing contexts is not at all clear and demands the development of new computational structures. In particular the work of Levy, Willsky and their students Tewfik, Nikoukhah, and Taylor has led both to novel structures involving both recursions that process data radially and parallel algorithms based on the spatial decomposition of the date to be processed. Also, the recent work on wavelets and multiresolution methods leads naturally to highly parallel algorithms with close structural and deep intellectual ties with the afore-mentioned multigrid algorithms.

Important contributions to global optimization using simulated anealing ideas were made by Mitter, Gelfand (in his doctoral dissertation), and Tsitsiklis.

In the meantime there was a growing realization that communication could be a bottleneck to the effectiveness of massively parallel computation. Based on available expertise in the communication area, a recent research effort by Bertsekas and Tsitsiklis and their students aims at developing efficient algorithms for handling the communication tasks that commonly arise in parallel computation. Some of the work of Tseng as well as the Ph.D. research of Luo on communication complexity was related to this effort.

The Laboratory's research on algorithms and distributed systems has brought a rapprochement between systems theory and theoretical computer science. Another point of rapprochement has been provided by complexity theory. As increased computing power became available, system theorists have been aiming at expanding the range of problems that can be solved numerically. While for some problems this was not too difficult, other problems were resisting. In research that started with the Ph.D. thesis of Tsitsiklis, and was continued jointly with Papadimitriou (then at Stanford), the tools of complexity theory were introduced to the systems field and were used to characterize the difficulty of the main problems of control theory, ranging from Markov decision theory, to decentralized control and to discrete event systems. In particular, it was demonstrated that complexity theory can be applied to control problems involving continuous (rather than discrete) data. Furthermore, computational intractability (in the sense of NP—completeness) was advanced as an explanation of some of the impasses reached in decentralized control and team decision theory, and provided an alternative to the more imprecise notion of analytical intractability. The Ph.D. research of Chow provided another contribution to this line of research.

Academic Visitors and Post-Doctoral Fellows

The research of the Laboratory in this period has been enriched by the presence of a large number of distinguished long-term academic visitors and post-doctoral fellows. We give an incomplete list: Luigi Ambrosio (Scuola Normale

Superiore, Pisa), Shankar Basu (Stevens Institute of Technology), Vivek Borkar (Tata Institute of Fundamental Research, India), George Cybenko (formerly at Tufts and now at University of Illinois), Bernard Delyon (INRIA, France), Peter Doerschuk (now at Purdue) Anthony Ephremides (University of Maryland), Bruce Hajek (University of Illinois), Hu-Bao Sheng (Republic of China), Jonnis Karatzas (Columbia University), P.R. Kumar (University of Illinois, Bernard Le Goff (INRIA, France), C.C. Li (University of Pittsburgh), Lennart Ljung (Linköping University), Yoshito Ohta (Osaka University), Kameshwar Poolla (University of Illinois), Anne Rougee (INRIA, France), Jawed Salehi (Bellcore), Les Servi (GTE), Jayant Shah (Northeastern University), Steven Shreve (Carnegie Mellon University), Alberto Sangiovanni-Vincentelli (University of California, Berkeley), Charles Van Loan (Cornell University), Gilead Tadmor (now at Northeastern University), M. Vidyasagar (University of Waterloo and Center for IA and Robotics, India), Eugene Wong (University of California, Berkeley), Ofer Zeitouni (Technion, Israel), Paul Tseng (now at University of Washington), and Denis Mustafa (MIT).

Chapter 4

The Center for Intelligent Control Systems

One of the most significant developments in the last few years has been the formation of the Center for Intelligent Control Systems, an inter-university (Brown, Harvard, MIT) center with its headquarters in the Laboratory. It was founded in late 1986 as part of the Department of Defense's University Research Initiative and supported by the Army Research Office with an initial five-year grant. Sanjoy K. Mitter serves as Director of the Center. Roger Brockett (Harvard University) and Donald McClure (Brown University) serve as Associate Directors. Several members of the Laboratory are members of the Center. From MIT there is also participation by Shafi Goldwasser, Silvio Micali, Ronald Rivest (all from the Laboratory for Computer Science), and B. Awerbuch, Thomas Leighton, Gilbert Strang and Daniel Stroock (all from the Mathematics Department).

An Agenda for Research: Towards a Framework for Intelligent Control

The fundamental contribution of classical feedback theory as developed by Nyquist, Bode and others was the demonstration that feedback around linear time-invariant single-input single-output systems can guarantee satisfactory performance even in the face of relatively poor knowledge of the system. While the initial theory was limited in scope, it established a mathematical context for the developments that followed. In time the state space theory of control allowed one to deal with time-varying and multivariable systems described by non-linear ordinary differential equations and in fact to specify control systems that optimize performance. Extensions to systems governed by partial differential equations and recent efforts to combine the desirable aspects of the original frequency domain theory and the theory of state space and optimum design have also met with success.

A crucial aspect of any successful control system is the extraction of information about the system to be controlled as the basis for determining the required control. Beginning with the work of Wiener and Kolmogoroff, and continuing with the notable contributions by Levinson, Kalman, and Bucy, a sophisticated theory has been developed for extracting the information from noise-corrupted measurements. Also, the theory of stochastic control for partially observed systems, spurred by Bellman's work on dynamic programming, provides a framework for combining information extraction and the design of algorithms that use this information as the basis for controlling systems.

The many developments and successes of the research efforts just described can be traced to several sources. One of these is that these investigations focused on well-defined classes of problems. That is, these theories are based on very specific classes of mathematical models of dynamics, uncertainty, observations and the information they contain, and desired objectives that allowed the development of a deep theory and effective analysis and design tools. Secondly, there were numerous applications to motivate and drive the theory and to benefit from the results.

In contrast to the theory of control as it presently stands, we have the field of artificial intelligence and related fields in computer science. The promise of AI has been to provide an extremely flexible framework for the development of intelligent systems for extracting information, decision-making, and control in situations in which our knowledge of a process is quite arbitrary in form and possibly non-numerical. Initial work in AI concentrated on the universal aspects of this problem with limited success. It was then recognized that it was not possible to develop methods for such an unstructured and general setting. Rather, knowledge of the specific process under consideration and of the task to be accomplished was essential. The concept of a knowledge-based system is perhaps the best-known example of this more focused approach to AI.

Donald E. McClure

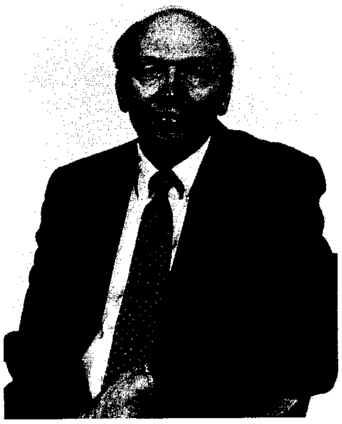

Roger W. Brockett

Sanjoy K. Mitter, David Mumford (Harvard), Tom Richardson and Gilbert Strang Discussing some mathematical problems in computer vision.

Thus at present we have the precise methodologies embodied in control and signal processing that have enjoyed success, focusing on highly constrained process descriptions but suffering from an apparent difficulty in adapting to the complexities that characterize certain advanced applications. At the same time artificial intelligence offers the promise of a flexible framework for tackling problems having much less structure but which seem to have begun to evolve in recognition of the fact that the use of structural descriptions of a problem is essential. In our opinion neither the theoretical framework of control and signal processing theory nor that of computer science and AI can by themselves provide the complete basis for a theory of intelligent control. For this reason, we have put together a research program blending the analytical and theoretical insights of control and signal processing with the problematique of artificial intelligence.

From one point of view, the main distinction between a theory of intelligent control systems and more traditional control theory lies in what might be called the "granularity" of the uncertain events with which the system is designed to deal. Traditional feedback theory deals with variables such as temperature and pressure which need to be regulated in the face of changes in the environment which affect their values. Intelligent controllers should deal with situations which may involve deciding what to control and the invention of strategies to achieve control in the face of altered configurations. Such a theory must go beyond the simple processing of signals or the estimation of the state of a system to the combined use of disparate sources of knowledge and data to extract the information required to make decisions intelligently. Such a theory must also be concerned with developing and exploiting structural representations of models at different levels of abstraction and must strive towards a mathematical codification of the concepts of learning, adaptation and organization.

A second aspect of intelligent control concerns the distributed and parallel nature of many practical systems. In dealing with such systems there is no alternative to the development of distributed algorithms for information processing, transfer, and decision-making. Of course, parallel systems allow one to propose and consider computational tasks that would have been discarded out of hand a few years ago. This in fact has created a new area of research, namely, the matching of computational and communication architectures and algorithms to problems. This area demands a theory of distributed computation, communication, information processing, and control.

In our research program we present a cohesive set of research directions that blend research in the classical domains of control and system theory, signal processing, computer science, and communication theory with the demands of new applications and technology to build on existing theories and to develop new ones that meet each of the challenges of intelligent control.[1]

The emerging field of intelligent control is of growing importance because of the burgeoning use of computer control in systems ranging from defense systems, automobiles and aircraft to inexpensive consumer electronics. It encompasses some of today's most exciting applications of technology. However, in spite of the abundance of examples, it is often surprisingly difficult to articulate governing principles having wide scope and significant predictive power or even to achieve agreement on a language for discussing the critical issues. This lack of a suitable codification has added greatly to the cost of producing new systems and prevented more effective use of some of its key ideas in fields such as manufacturing. In the Center, researchers balance their activities between

---

[1] Quoted from the Center brochure

investigations whose goals are to elucidate the fundamental principles of intelligent control which in some cases are already being used in applied work in an ad hoc fashion and activities which seek to establish a mathematical basis for discussing new ideas such as arise in distributed systems, learning algorithms, adaptation and control.

Although the work done in the Center is quite diverse, its main goals can be summarized as follows:

1. *To do research that will facilitate the design of synthetic intelligent control systems and contribute to our appreciation of naturally occurring intelligent control, for example, in biological systems.*

2. *To identify, and where necessary, create a body of literature defining the core science base of the subject of intelligent control.*

3. *To help establish the credibility of the subject through selected demonstration projects, done on campus or in cooperation with government (in particular, Army research laboratories) and industry at the forefront of technology.*

4. *To provide a clearinghouse for information on intelligent control.*

The research activities of the Center are organized into five areas:

Signal Processing, Image Analysis and Vision Control
Mathematical Foundations of Machine Intelligence
Distributed Information and Control Systems
Algorithms and Architectures

It should be emphasized that there is considerable overlap between these research areas.

To give a flavour of research in the Center we describe our efforts in the area of Learning. L. Valiant (Harvard) S. Mitter, R. Rivest, J. Tsitsiklis and their students are collaborating in this research.

Learning

The nervous system of humans and other organisms provide the inspiration for several of the major goals of computer science. Amongst these goals one of the most fundamental is that of creating computer systems that are able to learn from or adapt to their environment in the way even simple biological systems are able to do. In current computer systems essentially the only viable way of controlling behavior is to program them with every aspect of their expected performance detailed by a human agent.

Since software costs comprise such a major part of the cost of a computer system any significant learning capability would have evident economic impact. In addition, there are reasons to believe that learning is a fundamental capability that cannot be replaced by conventional programming in every instance. For example, large computer systems are created typically by large teams of programmers. When in use they often need updating. Even a minor modification to an otherwise correct system introduces the possibility of catastrophic errors unless the programmer involved has a full understanding of all relevant parts of a system. The advantage of a learning capability is that it would enable the update to be made with respect to the actual state of knowledge of the system, much as a human learns a new piece of information with respect to his current understanding of the world.

In everyday usage the world "learning" has numerous diverse meanings. In computer science one attempts to select some significant learning phenomenon, give it definition and try to understand it. One such phenomenon is that "of learning from examples," the ability to generalize from a few examples the intended concept, such as that of a "chair." There is little doubt that such phenomena exist. Children exhibit it to a remarkable extent. Experiments on animals such as pigeons suggest that this ability is not limited to humans. Furthermore, it can be demonstrated, if weakly, by computer experiments. The real question is whether we can understand this phenomenon sufficiently to be able to harness it in useful technology.

Substantial research is in progress aimed at understanding the ultimate possibilities and limitations of learning from examples, and of related phenomena where richer teacher-learner interactions are allowed. Finding a theoretical basis is especially important in this subject area. There is a substantial amount of experimental work over several decades that provides experience but no decisive conclusions. Also, in a system that learns we need to have some basis on which to trust its conclusions. We expect theory to provide a specification of what the system is accomplishing.

One unavoidable reference point for learning research is agreement on what models of learning or generalization really capture the desired phenomena. In recent years much progress has been made in this area. Having agreed on a model the next step is to select some styles of knowledge representation that are expressive enough to be useful and yet restricted enough that there is some chance of inducing them from informal interactions, such as the presentation of examples, rather than explicit programming. Such representations include Boolean expressions for functions, and finite automata for sequences. Also included are neural nets, which are inspired by biological nervous systems. As a result of current work an understanding is emerging about the possibilities and limitations of learning for such representations.

While our performance, as humans, in learning is spectacular compared with current computers, nevertheless we seem to learn slowly and with difficulty. It would be surprising if we could harness this subtle phenomenon for technological purposes without a careful understanding of it.

Charles Rockland

The Nematode: A Paradigm for Intergrative System Organization

The Nematode Project, initiated at the Laboratory by Charles Rockland and Sanjoy Mitter, is an attempt to elicit and develop the necessary theoretical tools by addressing the problems of integrative organization in the context of a rich yet comparatively simple biological model universe, the nematode. This organism provides both a focus and a rigorous testbed for theory development. The project is an inherently multidisciplinary effort, involving collaborators outside of MIT and coordination with both prior and ongoing biological experiments elsewhere.

The Nematode as a Model System

Nematodes, or roundworms, are from various perspectives the simplest multicellular organisms with differentiated tissue types. Over the past 25 years, largely through the influence of Sydney Brenner, *C. elegans*, a small (~ 1 mm long) transparent free-living soil nematode, has become the focus of intensive study from the standpoint of fundamental biology. The organism is amenable to a wide range of experimental techniques, which have yielded information spanning a multiplicity of organizational levels, from molecular to behavioral. The developmental lineages and cellular plan of the 959-celled organism are essentially invariant from individual to individual. A detailed "wiring diagram" has been worked out for the nervous system, which consists of 302 neurons, classified into 118 types. Complementary eletrophysiological data are obtainable from the related species *Ascaris*.

Despite the simplicity of its nervous system, the organism exhibits a variety of behaviors, some of which can be modified by training. This goes counter to prevalent views of nervous system organization, according to which complexity of behavior arises as an emergent or collective property of large numbers of (often homogeneous) elements. The nematode suggests a complementary viewpoint and direction of inquiry: How can organization capable of sustaining highly ramified and adaptive behavior arise out of the coupling and relations, or "constraints," among a possibly modest number of heterogeneous elements? A closely related theme is the non-modularity (or multiple modularities) of the organization.

Integrative Modeling of the Nematode

A central part of the project, serving as a bridge between the more theoretical work and the experimental biology, is a large-scale program of integrative computer modeling. The modeling framework, in keeping with the viewpoint noted above, is that of a family of "partial" models, of heterogeneous mathematical type, linked by interrelations or constraints. These partial models represent different components or facets of nematode organization, with the same object potentially represented in multiple models. A particular goal of the modeling is to determine the mutual-consistency constraints among the interacting control structures. Our expectation is that new mathematical structures will be needed here. We believe that two sources of relevant mathematical ideas will be: (1) Geometry and sheaf theory; (2) Theoretical computer science, and associated ideas of mathematical logic.

Postscript

Whither the Laboratory

In the previous chapters, we have described the evolution of the Laboratory over a fifty year period from the perspective of the conceptual and technological development of the broad field of Systems, Communication, and Control. Even a cursory reading of this document should convince the reader that today this field is inseparable from Computer Science. This does not mean that Communication and Control on the one hand and Computer Science on the other do not each have a separate internal logic of development. It does mean, however, that the conceptual overlap in the "problematique" of the areas is substantial and cannot, and should not, be ignored. In large part, this overlap arises from the enormous technological advances in electronic technology and the consequent trend to build larger and more complex systems. These large systems, whether their fundamental purpose is control, communication, or computation, tend to be distributed and to involve control, communication, and computation in a largely inseparable fashion. No adequate conceptual foundations exist for dealing with the inherent complexity of these systems, and this is recognized as the fundamental problem for the Laboratory in the future.

In the past, scientific progress on complex problems has generally involved the discovery of a methodological approach that made the complexity of the problem manageable. Such methodological breakthroughs usually involve individuals or collaborations of a few individuals, although there are often inputs from widely diverse sources. The puzzle for the future is whether these problems of large scale complexity will yield to the classical mode of research, or whether new modes of research are necessary. It will be seen from the current work at the Laboratory, and particularly the work at the Center for Intelligent Control Systems, that we are taking an intermediate position on this, continuing the classical mode of individual conceptual research, but utilizing inputs from much more diverse areas than in the past. How effective this will be, we do not know, but, with our control theoretic background, we will constantly use feedback from our success and that of others to adjust our approach.