

Academic Careers
for Experimental
Computer Scientists
and Engineers
.....

Committee on Academic Careers for
Experimental Computer Scientists
Computer Science and Telecommunications Board
Commission on Physical Sciences, Mathematics,
and Applications
National Research Council

National Academy Press
Washington, D.C. 1994

**COMMITTEE ON ACADEMIC CAREERS FOR
EXPERIMENTAL COMPUTER SCIENTISTS**

LAWRENCE SNYDER, University of Washington, *Chair*
FOREST BASKETT, Silicon Graphics Computer Systems and Stanford
University
MICHAEL M. CARROLL, Rice University
DON MICHAEL COLEMAN, Howard University
DEBORAH ESTRIN, University of Southern California
MERRICK L. FURST, Carnegie Mellon University
JOHN HENNESSY, Stanford University
H.T. KUNG, Harvard University
KURT MALY, Old Dominion University
BRIAN REID, Digital Equipment Corporation

Staff

MARJORY S. BLUMENTHAL, Director
HERBERT S. LIN, Senior Staff Officer
ARTHUR McCORD, Project Assistant (through February 1993)
LESLIE WADE, Project Assistant

**COMPUTER SCIENCE AND
TELECOMMUNICATIONS BOARD**

WILLIAM A. WULF, University of Virginia, *Chair*
RUZENA BAJCSY, University of Pennsylvania
JEFF DOZIER, University of California at Santa Barbara
DAVID J. FARBER, University of Pennsylvania
HENRY FUCHS, University of North Carolina
CHARLES GESCHKE, Adobe Systems Inc.
JAMES GRAY, Digital Equipment Corporation
JOHN L. HENNESSY, Stanford University
DEBORAH A. JOSEPH, University of Wisconsin
RICHARD M. KARP, University of California at Berkeley
KEN KENNEDY, Rice University
BUTLER W. LAMPSON, Digital Equipment Corporation
BARBARA H. LISKOV, Massachusetts Institute of Technology
ROBERT L. MARTIN, Bell Communications Research
DAVID G. MESSERSCHMITT, University of California at Berkeley
WILLIAM PRESS, Harvard College
CHARLES L. SEITZ, California Institute of Technology
EDWARD SHORTLIFFE, Stanford University School of Medicine
CASIMIR S. SKRZYPCZAK, NYNEX Corporation
LAWRENCE T. TESLER, Apple Computer Inc.
LESLIE L. VADASZ, Intel Corporation

MARJORY S. BLUMENTHAL, Director
HERBERT S. LIN, Senior Staff Officer
GREG MEDALIE, Staff Officer
JAMES MALLORY, Staff Officer
RENEE A. HAWKINS, Staff Associate
GLORIA BEMAH, Administrative Assistant
JANET QUARLES, Project Assistant
LESLIE WADE, Project Assistant

**COMMISSION ON PHYSICAL SCIENCES,
MATHEMATICS, AND APPLICATIONS**

RICHARD N. ZARE, Stanford University, *Chair*
RICHARD S. NICHOLSON, American Association for the
Advancement of Science, *Vice Chair*
STEPHEN L. ADLER, Institute for Advanced Study
JOHN A. ARMSTRONG, IBM Corporation (retired)
SYLVIA T. CEYER, Massachusetts Institute of Technology
AVNER FRIEDMAN, University of Minnesota
SUSAN L. GRAHAM, University of California at Berkeley
ROBERT J. HERMANN, United Technologies Corporation
HANS MARK, University of Texas at Austin
CLAIRE E. MAX, Lawrence Livermore National Laboratory
CHRISTOPHER F. MCKEE, University of California at Berkeley
JAMES W. MITCHELL, AT&T Bell Laboratories
JEROME SACKS, National Institute of Statistical Sciences
A. RICHARD SEEBASS III, University of Colorado
CHARLES P. SLICHTER, University of Illinois at Urbana-Champaign
ALVIN W. TRIVELPIECE, Oak Ridge National Laboratory

NORMAN METZGER, Executive Director

Executive Summary

Experimental computer science and engineering (ECSE) is the fundamental underpinning of the computer hardware and software that drive the information age. Many widely known computer advancements of the 1980s trace their origins to ECSE research. Examples are known to the average user—reduced instruction set computers, window systems, relational databases—but many more are hidden inside systems, making them faster, more efficient, or more functional, or they are part of the technological infrastructure that supports the rapid innovation so characteristic of the computer field. Experimental work is also an essential intellectual element of the computer science and engineering (CS&E) discipline that enriches research and teaching in the field.

Industry and universities both play major roles in supporting ECSE. ECSE research in academia in particular plays an important role in ensuring an adequate diversity of technical ideas out of which marketable products can emerge. Moreover, performing ECSE research in the university environment enriches the educational mission of universities and keeps faculty members current in the discipline. The importance of academic ECSE led the National Science Foundation (NSF) to request a study by the Computer Science and Telecommunications Board of academic career tracks for practitioners in the field. The NSF request was motivated by the concern that experimental computer scientists and engineers may face special chal

About this PDF file: This new digital representation of the original work has been recomposed from XML files created from the original paper book, not from the original typesetting files. Page breaks are true to the original; line lengths, word breaks, heading styles, and other typesetting-specific formatting, however, cannot be retained, and some typographic errors may have been accidentally inserted. Please use the print version of this publication as the authoritative version for attribution.

lenges in being evaluated and in creating appropriate research environments, although some top-ranked universities have met these challenges quite well.

THE NATURE OF EXPERIMENTAL COMPUTER SCIENCE AND ENGINEERING

Experimental computer science and engineering (ECSE) is a synthetic discipline in the sense that it studies phenomena that are entirely the product of human creation. Many interesting computational phenomena—processes, algorithms, or mechanisms that manipulate or transform information—are too complex to understand on the basis of direct analysis from first principles. For example, an algorithm may have many complicated states, or a process may involve time-dependent interactions of many subprocesses. For all practical purposes, such complex phenomena can be understood only on the basis of empirical observation.

Thus, ECSE refers to the creation of, or the experimentation with or on, computational *artifacts*. Often artifacts are hardware systems (such as computers) or software systems (such as text editors), but the term includes graphic images or animations, robots, or test and benchmark suites. When computational processes, algorithms, or mechanisms are implemented in an artifact, the behavior of the system and the interaction of its components can be observed in action. In general, when the computational phenomenon is complex, an artifact that embodies the idea will also be complex and will have many component parts.

In ECSE, artifacts may be the subject of a study, the apparatus for a study, or both. Artifacts often embody a substantial portion of the intellectual contribution of ECSE research, and their creation represents a significant intellectual effort.

Artifacts serve at least three primary purposes in ECSE. A given implementation can seek performance or seek improvement and enhancement of prior implementations (proof of performance), demonstrate that a particular configuration of ideas or an approach achieves its objectives (proof of concept), or demonstrate a fundamentally new computing phenomenon (proof of existence).

Computing artifacts are malleable and versatile. Unlike other machines, computers are "universal," meaning that within broad limits, whatever one machine can do, all machines can do. Although this property is extremely convenient in many respects, it implies the lack of an a priori limit on the functionality of computers, which feeds ever-expanding expectations for their capabilities.

The construction of computing artifacts is not strongly coupled to theoretical computer science. Unlike the more traditional sciences (e.g., physics) in which the interplay and coupling between experiment and theory are rather tight, an "experiment" in ECSE generally does not verify a prediction from theoretical computer science or rely heavily on a model developed theoretically. The reason is that the complexity of most real computing problems precludes the direct application of analysis: a problem can be made theoretically tractable only by abstracting so extensively that the problem that emerges may not capture the essence of the original problem.

Because ECSE is strongly coupled to technology, the availability of a given technology may well determine the feasibility of a good and innovative idea. The use of cutting-edge technology for a project subjects it to hazards such as instability, errors, and delays, whereas the use of a stable, mature technology carries the risk of obsolescence before the project is finished.

ECSE AND THE ACADEMIC ENVIRONMENT

The characteristics of ECSE have a major impact on how ECSE faculty must go about their work. An academic career in ECSE resembles one in other science or engineering disciplines only in outline (including teaching, advising students, writing papers and proposals, conducting research, serving on committees, and so on). The most substantive difference, perhaps, concerns how research is evaluated. Because ECSE is fundamentally a synthetic discipline, it is straight-forward to create a new computational phenomenon or an alternate implementation of a concept. Yet doing so does not automatically constitute an intellectual contribution. Rather, whatever has been created must be shown to be better than some alternative. In research in theoretical computer science, the key question is, Has the proposition been proved? In contrast, the key questions in evaluating ECSE research include, Does the idea provide a new and more useful capability or greater functionality? and, Is it faster or more efficient?

Box ES.1 describes how the mouse—a device for human-computer interaction—can be evaluated in these terms and thus found to be a success story in ECSE research. Other artifacts that emerge from ECSE research may not have had as illustrious a history, but they can also be significant.

Artifacts in a proof-of-performance role demonstrate that they are better because they are objectively faster or consume fewer resources. But it may be harder to prove the greater worth of artifacts

...serving a proof-of-concept or a proof-of-existence role, because the advancement may be qualitative, as in increased functionality. In this context, what is considered better may depend on subjective human judgments.

BOX ES.1 THE MOUSE AS ECSE RESEARCH

The computer mouse is used as a pointing device in human-computer interaction. The mouse was created by Douglas Engelbart at SRI as one of several human-computer communication devices. Although it was described in full technical detail and careful studies were made of its utility, many computer scientists recall first appreciating the power and significance of the invention, not from the published record, but from a film that Engelbart produced, showing the mouse in action. The demonstration of the device conveyed the essence of the new phenomenon beyond any amount of description of how, or how well, it worked.

In the ECSE research context, the concepts embodied in the mouse can be evaluated as follows:

- The mouse falls within the scope of ECSE, having mechanical, electronic, and software components concerned with human-computer interfaces.
- The mouse concepts fundamentally improve the functionality of the human-computer interface.
- The concepts were shown to be better quantitatively.
- The mouse has had significant impact, as witnessed by a variety of subsequent implementations, improvements, and applications, as well as its widespread use.

This lineage and history associate the mouse with a proof-of-existence role and place it squarely in the ranks of successful ECSE research.

The production of computing artifacts depends on a substantial infrastructure. Equipment needs are generally well understood, although few outside the computer fields realize just how quickly "state-of-the-art" equipment loses its cutting edge and utility for research. Equipment also requires space that must generally be specially equipped with power and air-conditioning capacity not found in standard office or teaching space. Experimental software systems often require dedicated or special-purpose hardware and cannot make use of the general-purpose computing environment that already exists in the department, school, or university.

Advanced graduate students are critical to ECSE research. Systems projects involve a great deal of detailed design and implementation, and because the artifacts involved in ECSE research must by definition be created from concepts (rather than blueprints), the students constructing them must have appropriate background, skills, and knowledge.

Larger ECSE projects require support from technicians and other staff who maintain the research environment and provide implementation assistance. Such staff free graduate students to focus more of their time on the intellectually significant parts of the implementation project and less on the more routine or lower-level (although necessary) components of system building.

Large- or even medium-scale ECSE research also requires collaborative or team effort, and all scales of research are likely to incorporate the previous efforts of others. Researchers can save valuable time and resources by using public domain and even commercial system components. More importantly, an individual researcher can have more impact by working with a team than would be possible by working alone and from scratch. Especially for large efforts, collaboration is essential because the subsystems of an artifact are often so specialized that other expertise is needed.

Funding, which is important to researchers in any field, must be sufficient to cover the long time horizons and large demands on resources that characterize ECSE research. However, programs intended to support junior faculty who have not yet established their professional reputations (e.g., the NSF Young Investigator program and the Research Initiation Awards program of the Computer and Information Science and Engineering Directorate of NSF) are highly competitive, and so only a few faculty members receive such support in their first year after graduation. Thus, new ECSE faculty are not likely to have research support based on their own research ideas during the early years of their careers and may well have to rely during this time on "start-up" funding provided by the hiring school.

Finally, the time demands on ECSE faculty are inflated by the nature of the field. For example, given the time needed to build a research team and a laboratory, even a very good assistant professor in ECSE may have only one completed Ph.D. student and/or one major completed project before a tenure decision is made. Because the dissemination of artifacts is an essential medium through which research knowledge in ECSE diffuses into the community, the six-year probationary period for an assistant professor may be too short to establish a reputation in the field.

ECSE researchers share their software, provide access (via Internet)

to their experimental computers, and distribute their data files. Disseminating by means of artifacts allows research colleagues to acquire a working knowledge of an experimentalist's accomplishments that is deeper and more extensive than would be possible simply by reading journal papers. In some cases (e.g., in the display of a graphic animation that one must see develop in time), there is no adequate written alternative to observing the animation in action.

The ECSE research community depends heavily on conferences to communicate new knowledge, and conferences are widely regarded as the preferred medium for maximizing the intellectual impact of ECSE research. However, the tenure and promotion process at many universities does not give conference presentations and publications a weight appropriate to their significance in the field, preferring instead publications in archival journals.

The focus of ECSE on artifacts has often led to a tension between theoretical and experimental computer scientists. Although the tension seems not to manifest itself in some departments, in others it appears to have caused experimental or theoretical work to be misunderstood and underappreciated by researchers in the opposite camp. This has obvious implications for junior faculty members in departments in which the senior faculty are primarily of a "different stripe"; their readiness for promotion and tenure may not be evaluated according to the standards for quality and the criteria for success that apply to their research specialty.

PROVIDING A NURTURING ENVIRONMENT FOR ECSE FACULTY

Some schools, typically those with large groups of experimentalists, have fostered highly supportive environments for ECSE faculty. However, many more schools—perhaps because of their smaller size or particular history—have few experimentalists and little experimental activity under way. In such schools, many ECSE faculty perceive the career environment to be difficult or hostile.

One key element of a positive environment for ECSE is the availability of a mentor for junior faculty. Mentors should provide advice on issues such as publication, funding sources, collaboration, choice of problems, and logistics. For example, a mentor might suggest conferences or journals in which publication would bring maximum exposure or prestige, funding agencies most likely to support a junior faculty member's work, senior professors at other institutions with whom to collaborate, and ways to structure projects so that intermediate results could be made available.

Advocacy of a faculty member's interest to higher authorities (e.g., university administrations) is also essential. Advocates would, for example, argue the case for obtaining review letters from appropriate parties, explain key characteristics of ECSE to those outside the field, accumulate evidence of the impact of the junior faculty member's work, and document the structure and stature of the literature (e.g., which conferences and journals are respected, prestigious, and well refereed).

In addition, department and university evaluators must strive to use standards and criteria for tenure and promotion decisions that normally characterize productive work in the ECSE discipline, rather than standards that may be applicable to more traditional academic disciplines, taking care not to exclude meaningful evidence of achievement (e.g., artifacts with substantial impact on the ECSE community) simply because it is nonstandard.

With respect to the important letter-writing process in tenure and promotion, the primary criteria in selecting potential letter writers should be their stature in the field and their familiarity with the candidate's work. Factors such as the letter writer's institutional location or status as a collaborator should not be reasons for excluding letters. In particular, because views from industry may be important for judging the impact of ECSE work, letters from individuals in industry or government laboratories should not be arbitrarily limited, and they should carry equal weight to those of similarly qualified and reputable individuals in academia. Similarly, eschewing letters from collaborators in a field as intrinsically collaborative as ECSE is to eliminate some of the best possible input regarding a candidate's intellectual capability, creativity, and originality.

More generally, universities should recognize that an experimentalist being considered for tenure or promotion may have fewer publications (and predominantly conference publications), nonstandard forms of dissemination (e.g., distribution of software artifacts), substantial amounts of collaborative research, and few graduate students completed, and yet still be a spectacular researcher. A judgment should be based on the presence or absence of the following:

- One or more computational impact-producing artifacts completed;
- Research results disseminated to and used by the community;
- A reputation for novel systems solutions or ingenious experiments; and
- A filled or filling pipeline of well-trained graduate students.

It is the responsibility of the candidate to achieve distinction. It is the responsibility of the department and institution to recognize and reward it.

Departments can improve the environment for ECSE faculty by providing technical staff support and laboratory space. Start-up packages for new assistant professors in ECSE comparable to those received by new experimentalists in other departments would enable them to begin research more quickly. Providing opportunities for junior faculty members to teach advanced seminars in which graduate students can receive needed training in preparation for joining a research project would facilitate the building of a research team.

The federal government and industry also have critical roles to play in improving the environment for ECSE in academia. Most importantly, the federal government should realize that a variety of funding structures are needed to support ECSE research, including small, relatively short term grants or contracts that focus primarily on the development of a concept; medium-scale group funding; and large, relatively long term grants or contracts associated with deliverable computing artifacts. Postdoctoral support for new Ph.D.s in ECSE would help to overcome some of the limitations and constraints imposed by the six-year probationary period for assistant professors.

The computer industry can help to enhance the environment for ECSE in academia by establishing collaborative work arrangements with universities, including those that may not be nationally known or recognized. Computer or software companies that interact with such universities and thus expose local CS&E departments to the problems and needs of industry not only foster meaningful collaborative work but also help to produce students who are better informed about these problems. Such students graduating from less well recognized universities may be more likely to work for local computer or software companies. If appropriate nondisclosure agreements can be achieved, industry can also provide access to hardware designs or source codes for various software systems. An academic researcher's access to source code will certainly reduce the time required for him or her to complete an experimental software system and may result in an improved system of direct interest to the owner of the source code.

About this PDF file: This new digital representation of the original work has been recomposed from XML files created from the original paper book, not from the original typesetting files. Page breaks are true to the original; line lengths, word breaks, heading styles, and other typesetting-specific formatting, however, cannot be retained, and some typographic errors may have been accidentally inserted. Please use the print version of this publication as the authoritative version for attribution.