Dick Mackinnon

HISTORIC

~~IBM CONFIDENTIAL~~

General Observations on The Role, Value and Problems of The Scientific Centers

by
Harwood G. Kolsky
IBM Fellow
November 26, 1973

I. Introduction

Since the beginning of the Scientific Centers in 1964, management has experimented with many procedures and organizations concerning the handling of the individual projects, the staff members and the facilities. Although it is not yet clear that the best organization and procedures have been achieved, many of them have stood the test of time and have proved useful in practice.

The World Trade scientific centers, first started in 1968, have been building steadily following essentially the same principles and mission statements as the domestic centers. Although World Trade has many problems unique to the individual countries in which the centers are located, the basic problems of the relationships among the scientists, the centers and IBM management are very similar.

During the last three years the World Trade centers have been prospering and growing while the domestic centers have been shrinking under mounting pressure. The reasons for this are very complex, but they certainly are not because the domestic centers have failed in their mission nor because IBM is disappointed in their accomplishments.

Much of the material in this memo has been presented before in various documents, but it is still as valid as when it was first written. I only hope that the arguments can help turn the tide and start a renewed strengthening of the centers which represent an extremely productive but rare and fragile resource to IBM.

II. Mission and Objectives of the Scientific Centers

The mission and objectives for the Scientific Centers has been well stated in the long range plan for the Data Processing Division originally prepared in April, 1968 (see the Appendix). The environment of the Scientific Centers may be described as a world, which is rapidly changing, based on increasingly difficult technologies and increasingly complex problems of our modern urban society. IBM in the past has owed much of its success to its ability to utilize advanced scientific concepts in practical applications. At the present time our share of the scientific market has been decreasing steadily under the pressure of competition. The traditional IBM methods of developing and marketing products have not been as successful in this area of the business as they have in the more commercially oriented fields.

The Data Processing Division has felt a considerable need to expand the computer market by solving customers' problems in totally new areas of endeavor. The industry development activities work mainly on implementation of existing application technology on a short time scale. Between these activities and basic research as done in the Research Division there lies a considerable gap. The Scientific Centers were originally instituted to help fill this gap by permiting direct contact with advanced scientific institutions, both universities and customers, in the localities in which they are located. The Centers attempt to prevent the erosion of IBM's position in the scientific market place by building the image of the corporation as a company interested in scientific applications. They also promote the direct sale of equipment by identifying and initiating new areas of computer applications and, on occasion, by direct sales assistance on scientific benchmark programs, etc.

III. The Effectiveness of The Scientific Centers

The DP Scientific Centers serve two IBM needs in the real world of long range customer requirements: application research and scientific image building. In the opinion of independent auditors the Scientific Centers serve these needs better than any other IBM organization. The reason for this is that the Scientific Centers are close enough to the real world to feel the pulse of customer requirements and to properly weigh technological trends that will influence DP sales. Yet they are sufficiently detached from the marketing organization so that the academic and scientific communities treat them as peers rather than salesmen. The need the Scientific Centers fill is a DP need since it relates directly to market definition and market development.

In applications research, the Scientific Centers identify new computer applications and demonstrate their feasibility at low cost. To a degree unique in IBM, The Scientific Centers have been able to transfer their results to DP in the form of programs and RPQs. This success is due to being part of DP -- in day to day contact with DP headquarters and branch office people -- and to the fact that most of their projects are carried out with customers who don't hesitate to bring pressure on DP to implement new ideas.

In scientific image-building, the Scientific Centers maintain professional level contacts with university and industry scientists who influence the specifications for computer procurements. These people have generally negative attitudes toward IBM which, to them, is a manufacturing firm with no interest in science for its own sake. The Scientific Centers reduce this negativism by demonstrating that we do have bright people working in challenging areas relevant to the interests of the scientific community. They also increase customer receptivity for new applications. Scientific image-building requires effective university contacts. The Scientific Centers have achieved this at certain schools. It should be noted that these contacts are with university professors, department heads, and graduate students. The fact that IBM executives and salesmen deal with different people at the university frequently masks their appreciation of the Scientific Center presence on campus.

The overall quality of Scientific Center work is high. The productivity as measured by publications, application programs, external contacts and demonstrations is very high compared to other groups of similar size. The quality of Scientific Center work depends on having a critical mass of good people in several disciplines. The inter-disciplinary mix leads to maximum cross-industry activity. Small groups are forced to work with external people to get things done -- this is advantageous in the Scientific Center; however, the group must start with a viable critical mass. Too large a group should be avoided because it creates undesirable project-oriented nuclei and reduces Scientific Center flexibility. The critical mass (of 25 - 30 total staff plus computer) allows meaningful work to be carried beyond the "idea" or "survey" stage.

IV. Questions Facing The Scientific Centers

Whenever Scientific Center managers gather, certain problems are mentioned again and again as being important to the health of the Scientific Centers. Most of these problems are not the kind that can be "solved" once and for all by some management action. They are more like constraints which must be continually considered in all decisions. The way in which management copes with these questions on a day-to-day basis will set the course of the Scientific Centers during the coming decade.

A. Questions of Mission

The domestic Scientific Centers were established originally in 1964 and the WTC Centers started in 1968-69 to fill the needs of their respective countries in the gap between basic research and technology on one hand, and marketing development activities on the other. There is also the increasingly important mission of maintaining direct contact with the advanced technological institutions and Government agencies, with the resulting impact upon IBM's scientific image in each country.

Because of the three important aspects of the mission--research, marketing, and image--the Scientific Centers are always being pulled toward one or the other by management. The most successful projects have a proper balance between them.

The concept that the Scientific Center manager should have an entrepreneural role in selecting things to do which will best serve the Center, science, and IBM's image and market, is not really understood very well even by its own management.

The ability to see an opportunity which matches the skills of the Center and to move fast to do something about it in a hurry is a great strength of the Scientific Centers. The nomination and approval cycle for usual projects in IBM prevents us from acting as quickly as competitors. The Scientific Centers can act as "internal competitors" for IBM and cut through technical difficulties.

The Scientific Centers give to IBM some of the advantages of a small company. Having small groups of sharp people, having long experience in their respective fields, running hard and fast on well-understood projects can result in incredible achievements. One could argue for having a few refuges within IBM for such activities on purely philosophical grounds, even with no immediate return on investment. The fact is that the return is better on a per-capita basis on almost any measurement which can be applied.

The ability for Scientific Centers to use new equipment and exploit it gives valuable early feedback to SDD. The ability to do things for a small market or an unproved market means high risk projects can be started.

The Watch Dog aspect of the Scientific Centers should not be overlooked. It is much harder for other IBM groups to do extremely poor jobs or outrageously expensive jobs when faced with Scientific Center evidence or statements to the contrary. This same philosophy is followed by the Research Division in its areas of scientific competence.

B. The Reporting Position Question

Because the mission of the Scientific Centers is balanced between research, market, and image, it is possible to have them report to any organization having one or the other of these missions in its charter. It can also report either at a high or low level in such an organization, with either a line or staff function.

The domestic Scientific Centers began reporting to Corporate Headquarters in a research, staff-like position. They were later shifted to DP, mainly for financial reasons. Their reporting point within DP has shifted several times over the years, indicating that there is uncertainty as to the "correct" position.

The WTC Centers report either to DP or External Affairs (market or image) within their countries, although there is considerable variation in the exact conditions. The recent trend has been toward External Affairs. This staff position makes it easier to operate at a policy level on scientific affairs, since External

Affairs acts as an arm of the country general manager. It also makes entering into joint studies, fellowships, etc., easier because the approval cycle is short and contacts are made at high levels in customer organizations. The main problem with such a corporate staff position is financial support. DP has a much larger budget than any other organization, so the Scientific Centers will be dependent upon it indirectly, if not directly. Even if The Centers reported elsewhere, they would have to maintain good relations with DP on individual projects.

There is probably no single position which is "right" if by that one means no conflict will arise. The reporting point should be re-examined periodically to see if a change would help. In the final analysis, The Centers should report to that senior executive who is best willing to support and utilize them as Scientific Centers, not as something else.

C. The Question of How to Allocate Resources.

Financial problems always rate high on a Scientific Manager's list. Financial support is a concrete expression of the belief of IBM in a particular organization. It is interesting to ask how discretionary resources are assigned within IBM and match these against the Scientific Centers' needs:

(1) The first is the formal way, involves making marketing forecasts and building a business case before starting serious work on a project. This method works well on large projects which are extensions of IBM's existing product line. It does not work well on new opportunities, e.g., graphics, terminals, pocket calculators, or copiers. It also does not work well on small products with a limited market where conservative pricing policies can be deadly.

(2) Another approach is to find a project based on the inspiration of a charismatic leader. This works best on new opportunities previously overlooked by the company. The most spectacular advances come from such projects, but they can result in spectacular failures as well. The failures tend to generate more controls and more bureaucracy for future project leaders to face.

(3) One approach frequently used is that of estimating the return on investment while comparing two or three similar competing projects. Too often the resources go to the group which can put up the best P.R. story. The management problem is one of comparing unlike things. A "paper tiger" is always better than an existing entity whether it is a hardware product, a software product, or the plans for a research project.

(4) Sometimes resources are applied on the basis of past performance. Intuitively one feels the best ways for IBM to assign resources would be to reward the people or groups who have a good record of accomplishment and punish those who have poor record. All too often those who have squandered resources, missed milestones or produced shoddy products are rewarded by being given more resources to catch up, while those who are frugal and efficient are cut back.

The Scientific Centers would certainly prosper if more decisions were made on the latter basis.

D. The Question of the Optimal Size for a Center.

A frequently asked question is, "What is the optimal size (or critical size) for a Scientific Center?" I usually answered the question by giving the following three experimental points:

(1) The Los Angeles Scientific Center had about 75 or 80 people when the Scientific Centers were formed in 1964. It was gradually worn down as being too big. Both the Palo Alto and Houston Centers were really built from its headcount.

(2) The Washington Scientific Center was formed in 1966 without proper funding or management support. It grew slowly to about 12, was stationary, and finally was recombined into the other Centers because it was too small.

(3) The Centers which have survived and thrived have been staffed with 25-35 people.

Based on these three experimental points, one can say that 75 is too large, 12 is too small, and 25-35 is good. The really "optimal" size depends very much on the individual location, projects, personnel mix, etc. The closeness of coupling to neighboring universities and other scientific groups is particularly important.

E. Personnel Questions

There are two chronic personnel problems in Scientific Centers: (1) The hiring and retention of good staff members, and (2) the career path problem.

Standards for excellence may differ from one professional activity to another, but as far as Scientific Centers are concerned, there should be no compromise with quality. A Scientific Center does not deserve the name if it tolerates poor work on any of its projects. The way to ensure technical excellence is, of course, to have excellent people. One of the main tasks of the Scientific Center manager is to select and properly motivate good staff members. Excellent employees are easy to manage; those of average ability take more guidance to make sure that they do good work. Poor staff members should be eliminated.

Because of the mission, Scientific Centers need staff members who (1) are good practicing scientists, (2) have a sense of the marketplace, and (3) have good judgment in dealing with other scientists, customers, and IBM management.

Having hired or developed these rare scientist-marketeer-statesmen, the problem inevitably arises as to what the career path is for such individuals. The same question occurs for scientists anywhere in IBM, but because of the small size of the Scientific Centers, it is particularly acute. This is clearly a continuing problem which must be faced on an individual basis. Programs such as rotational assignments to other Centers or to other organizations within IBM, sabbaticals to universities, etc., can help.

The Scientific Centers also participate in the WTC Fellowship Program. There is a feeling that returning scientists can better identify with the local IBM Centers as a result.

F. The Questions of Project Selection and Monitoring

One of the most commonly asked questions is, "How are Scientific Center projects selected, monitored and killed?" A lot of time has been spent both in certain Scientific Centers and at headquarters explaining the procedures to be followed. Flow charts with several classes of criteria, multiple decision points, weighting factors, etc. are sometimes shown. Procedures are described which were capable of screening dozens of projects per week. Yet is is also obvious when discussing the actual projects being done that they are not done this way. In a given Center, old projects are killed and new projects are started only two or three times a year, so each is really studied as a special case in practice.

It has been observed in the Scientific Centers for some time that the really good projects come from the efforts of individual staff members who push ideas to fill a need which they can see because of their intimate knowledge of a field. Very few projects are ever handed down from above, or even suggested from Industry Development. Certainly none ever came through a formal flow chart of "criteria" and "decision points."

Once a project is underway, the question becomes one of monitoring it in a regular way. Formal project descriptions and regular progress reports are, of course, standard. Technical reviews are discussed under "measurement."

The grid of project highlights (see example in The Appendix) as used for several years, was a very useful way to survey projects. If the grid is accurately filled in, it automatically answers most of the simple questions people will ask about a project. The grid is also a convenient way of summarizing information which can be put on flip charts at a moment's notice for presentations.

G. The Question of Terminating Projects

The problem of terminating a project which is going nowhere is one of the more difficult management problems in the scientific environment. The arbitrary or capricious starting and stopping of projects will create lasting resentment and stifle individual initiative very quickly.

Too often the problem is approached from the point of view: "This project shall be killed to save resources, headcount and dollars." A much better way would be to ask,"What can we do which is more valuable with these particular people, with these particular skills, at this particular location?"

The best way to terminate a project is to start a better one. The people concerned will then switch over willingly rather than fighting desperately to keep their old project alive.

Stopping a well-run project simply because it has no short-term payoff is certainly contrary to the Scientific Center mission. For example, the plasma physics project at Palo Alto was judged in September of this year as a prime candidate for cancellation because it was "irrelevant to FS." Now, a month and a half later, because of the energy crisis, it is one of the hottest projects in any of the centers. Fortunately, it was not terminated in the meantime.

H. The Questions of Measurement

One of the oldest questions relating to scientists in industry is, "How can one measure the productivity of a scientist, a project, or a Center?" In the DP organization where everything tends to be reduced to points, this is particularly important.

If simple, definitive measures can be made, then decisions become simple and definitive. Unfortunately, scientific work and modern computer systems work is complicated and involved when it is new and important. It tends to be a personal sort of thing, where only certain people can make certain things happen.

The secret certainly is people in the Scientific Centers. A Scientific Center peoject usually depends on one or two key individuals. This Technical leader knows better than anyone else what the difficulties and the promises of the project are. As with our scientific customers, his judgment must be one of the key deciding factors in what the project is and how well it is doing.

To aid in the resource allocation problem and to avoid the problem of blind spots in projects, it is important to have technical and marketing reviews of the larger projects on a regular basis. The Technical Advisory Committee and the Marketing Advisory Committee were used for this purpose for several years as a management tool by the manager of The Scientific Centers. If done properly, these reviews are more constructive than external "audits," which usually imply there is trouble with the project.

The method of these reviews is the same as the scientific method--the submission of one's ideas to the scrutiny of his peers. This is the main way scientific progress is made and ideas propagated.

We should realize that, aside from salesmen on quota, few groups in IBM have accurate measures of their activities. Those calling for better measurements of scientific projects, where it is admittedly hard, really do not have them on development nor service functions either. They certainly do not have good measurements on the quality of forecasts, industry analysis, financial or personnel planning, etc.

I. The Questions of Communication and Duplication

The bugaboo of "duplication of effort" ranks high in upper management's worries about projects in IBM. In The Scientific Centers it is hardly ever a real problem. Even projects having the same name are usually quite different in content or approach. As Manny Piore once said, "Duplication of projects is not necessarily bad. It is the lack of communication between similar projects which is unforgivable."

The lack of communication between related projects is a common problem throughout IBM. The Scientific Center Managers meet on a regular basis four times a year. In addition, joint seminars with other organizations are sometimes held. However, although the centers receive each other's completed reports, there is not enough informal interaction on work-in-progress. Joint projects between the Centers are frequently suggested but rarely done.

The best solution is to have staff members visit other Centers and give lectures on their work. Reports without personal contact tend not to be read or understood.

J. The Transfer Problem

The transfer of the results of scientific projects to other groups within IBM, or directly to customers, has been a continuing problem throughout the years. There is little I can add to this topic that has not been said dozens of times. A project leader normally has to spend an appreciable part of his personal effort in seeking customers for his products. The percentage of success still remains questionable, since it is frequently hard to tell when a scientific "product" has been transferred. A concept used by another group can be very valuable although no paper transfer shows. Our successes in the systems area are easier to document.

K. The Approval Procedure Problem

Each of the Scientific Center managers will sooner or later mention the difficulty of getting approvals. Examples include new projects, joint studies, fellowships, special equipment, etc. Usually the difficulties were separated from purely financial questions.

Scientific Centers tend to do unusual things, with the result that whatever they propose is likely to be a special case and tends to get stuck in the decision-making gears at all levels.

Unlike some of the other problems on this list, this one does have some hope of solution as management and staff people become educated to the importance of Scientific Center projects.

APPENDIX:

From "Scientific Centers Long Range Plan" apr 23, 1968

MISSION

The Scientific Center mission is fourfold:

Communication: A vital part of Scientific Center mission is to maintain an effective link between various technological entities both within and outside IBM. Such entities would include Industry Development and Marketing, Universities, Advanced Technological Customers, Product Marketing, IBM software and hardware design group, and other IBM Research and Advanced Technology activities.

Initiation: Scientific Centers are also directed towards identifying advanced technological application and system areas whose development would broaden the computer marketplace. Scientific Centers would initiate projects and execute feasibility studies to give initial impetus to these new areas of endeavor.

Propagation: Scientific Centers will work, primarily with Industry Development groups, to provide transition of advanced development activities and technical consultation. In order to complement this aspect of the mission, Scientific Centers will also work, primarily with SDD, to assist in initiating appropriate hardware and software development activities necessary to maintain and increase IBM's position in the marketplace of the future. They will also propagate their activities throught education, demonstration and other modes of dissemination.

Image: Scientific Centers will maintain high image in the outside technological community in order to be focal points for attracting high-caliber technologists to DP Division and the IBM Corporation as a whole.

[redacted]

Good general Statement

STRATEGIES

I. <u>Establishment of Scientific Centers</u> In order to fulfill Scientific Center objectives, such Centers will or have been established in seven key locations. The factors that influence the choice of locations (Development Center locations, Population, Industrial Activity, University environment, etc.) are discussed in Appendix A. These locations are Cambridge, Chicago, Houston, Los Angeles, New York, San Francisco Bay Area, and Washington, D.C. We feel that there is adequate national coverage with the establishment of these seven Centers, and that it is preferable to allow these Centers to grow to full viability and stability rather than to embark upon establishing more Centers of this type. However, assuming normal growth, we can foresee a time in the future (five to seven years) when the establishment of further Centers would be required to keep pace with, and to ensure support of, expanding Industry Development activities.

II. <u>Control Size of Scientific Centers</u> The individual Scientific Centers must be of some reasonable size. We feel that a twenty man Center, achieved over a two years span, is critical minimum size. However, a viable stable size is about twenty-five people on the following basis:

- It is enough to be seen and respected in a given community
- It provides a reasonable balance between productive scientists, programmers, managers, and secretaries.

- Work based on about three or four disciplines can be carried on with each scientist having co-workers with whom to discuss problems.
- Work can proceed simultaneously on four or five projects, thus avoiding the inherent instability of a single project organization, and yet ensuring the dynamicism necessary to avoid over-specialization.

Probably thirty is a maximum useful size; beyond this the Center would start to get too introverted and over-organized. This number of from twenty to thirty assumes that other supporting services are proved, and that the total Center activities on site, including the Scientific Center and Industry Development groups, is of sufficient size to warrant a substantial machine configuration consistent with needs.

In discussing the size of a Center it is important to remember that scientific talent cannot be rapidly be shifted from a project to a radically different one. In spite of the fact that we attempt to recruit personnel who are not specialized to a degree beyond that necessary to attain a high level of professional competence, re-orientation of scientific personnel has to be exercised with a great deal of care and good judgement. In particular, if the size of a Center is to be changed, considerable lead time is required since top level scientific talent is not easy to recruit or, if found to be "surplus", placed in other IBM positions. It is thus important to have long term plan as to the size of each Center.

Important

V. <u>Definc Work Areas</u> It is intended that projects be initiated at the Scientific Center level, subject to the requirements laid forth in the devinition of OBjectives. These guidelines are established to ensure that such projects are meaningful to the over-all Scientific Center mission, and which can be used as a basis for management evaluation and decision, and include the following aims:

- the project should lie in a field which is abreast of the frontiers of science and technology, thereby being a focal point of intellectual excitement
- the project should lie in an area which is potentially of substantial commercial value to the DP Division of IBM, even though we recognize that it may not be practicable to establish a firm market projection in terms of NSRI points until the project begins to approach its developmental phase.
- the project should be such that due to its high level of interdisciplinary and/or professional requirements, it cannot be most effectively executed by other DP development groups
- the project should be one that can benefit most (in terms of speed of progress) by the particular Scientific Center environment, especially in terms of the communication links established between disciplines, with other DP groups, other IBM Divisions, and, in particular, outside technological institutions and other customers.

HIGHLIGHTS OF PROJECT: ATMOSPHERIC PHYSICS

	ENVIRONMENT	OBJECTIVES	STRATEGIES	RISKS, CONTINGENCIE ACTION PLANS
ENTIFIC	1) Numerical weather calculations have been one of the traditional areas used to justify super computers 2) World Weather Programs are now in an active planning phase.	1) Complete and publish significant scientific research in the field. 2) Influence IBM's large computer projects.	1) Select key technical problems (e.g. radiation transfer) and work on them in depth. 2) Develop program for radiation scattering in atmosphere July 1969. 3) Monitor scientific developments in related fields.	1) Long-term continuity of effort is required to make meaningful scientific cor butions to the field. 2) ACTION--recognition of this fact within DP.
RKETING	1) Field is heavily dependent on federal funding. 2) IBM has lost over half the market since 1964. 3) Super computers are now being developed by IBM and by our competitors.	1) Build image and provide support to the field in helping regain the environmental science market. 2) Participate in various environmental science programs. 3) Prepare strategy papers for World Weather, etc.	1) Concentrate on personal contacts with key leaders in the field. 2) Work with GEM, WTC and FSD representatives serving the field.	1) IBM will market a super computer in time for the early 1970 procurements. ACTION--Emphasize the importance of the scienti market.
ANCIAL	1) Project is part of the Applied Physics project which is funded for 1968-1969.	1) Not to exceed the Scientific Center and Applied Physics budget.	1) Maintain at least a level of support sufficient to ensure continuity of effort.	1) Funding for the US and W Weather programs become available on schedule. 2) Other applied physics projects may preempt all of resources. ACTION--Seek additional support.

10/21/68
HGK

MANAGING an ADVANCED TECHNOLOGY LABORATORY

Draft for "Sloan Management Review" Article
July 18, 1988

INTRODUCTION

Annually the IBM Corporation spends a considerable amount of money for research and development. In 1984 this figure was reported as XXX million. Its research activities are concentrated in the Research Division, headquartered in Yorktown Heights, New York while development and engineering activities are spread throughout many product and marketing divisions worldwide. Organizationally, IBM looks similar to many other industrial corporations, especially in its distinction between product development and research organizations.

However, within IBM another technical component can be said to represent a non-traditional approach to technology. Such are the scientific centers which have existed within IBM since the early 1960s; they perform an unusual and unique role within the company. Their management and technical staff view themselves as the intrapreneurs of modern terminology and draw upon an experience and heritage of nearly twenty-five years -- corresponding to the adolescent period of computing.

This article explores the nature of this IBM organization and discusses how the technical agenda is managed. A by-product for the reader will be insight into their operation and contributions. The scientific centers exist within a corporation which is well-known for its willingness to pursue technical ends in a highly-structured fashion with significant resource commitment. Somewhat irreverently, this has been referred to by some observers as the "cast of thousands". As we shall see, the scientific centers undertake their work in a quite different mode characterized often as counter-culture and is, at the very least, different from the rest of the corporation. Therein, as we shall see, lies their strength of contribution as well as the seed of their long-term frustration. A delicate balancing act by center management is the result.

IBM scientific centers exist as three laboratories within the United States (in Cambridge, Palo Alto, and Los Angeles) and in some 15 countries within IBM's World Trade Corporation. This article focuses upon the domestic activity since the World Trade centers have a somewhat different objective. They are aligned more closely with the particular interests of the country in which they reside. Also, the discussion will focus upon the Cambridge Scientific Center, given the author's responsibilities and the identical charter enjoyed by all the United States scientific centers.

MISSION and ORGANIZATION SETTING

Cambridge Scientific Center's mission is to undertake advanced technology work, primarily in systems and programming. Universities and the academic community more frequently call such work **applied research.**

The scientific centers are one of the few places within IBM whose **sole** function is to pursue adtech activities. A 1977 study conducted within the company showed fewer than five hundred people worldwide engaged in such work. Since Cambridge has a total population today of forty-three (including administrative and managerial), a key theme or characteristic is that such operations are small and rarely live beyond a critical mass level of staffing. Hence they are also frequently obscure and thought of as "closet" operations within IBM at large. As I hope to demonstrate this becomes a two-edged sword and not the least of the challenges in their management.

Defining adtech activities has historically been both a necessity and a problem as scientific centers are portrayed to visitors and inquisitive management executives. My definition is hardly rigorous or precise but says that adtech work in IBM lies between the two extremes of research and product development activities. In short, we generally do not set out to engage in research or to deliver a product but, instead, pursue another agenda and march to another proverbial drummer. This can be characterized as experimental work which is quite applied in nature and often is aimed at proving the feasibility of a concept, idea, or approach. We may know what we want to do but how to do it and how well it can be done may be the pressing issues.

The horizon for our work is normally 3-4 years and at any particular time certain projects may indeed appear to have research or product emphasis and flavor, but as we shall see, this relates more to the evolution or cycle of the project than anything.

Organizationally, the center is part of Academic Information Systems (ACIS), a new special business unit formed in 1982 and which focuses upon campus computing in the academic and research arenas. Prior to this the scientific centers (since 1964) have always reported to one of IBM's line marketing divisions -- initially the Dataprocessing then its successor National Accounts Division. Today ACIS has largely a marketing flavor and orientation. So one concludes that scientific centers have always been administratively aligned with IBM marketing.

In my view this arrangement is one of the significant decisions made at the birth of the centers -- and reinforced at crucial times along the way. It colors our work and is not insignificant in accounting for our successes and the way we play our end-game. For we feel we are very much involved in problems of the real world and impose considerable pressure upon ourselves to be deemed relevant to our reporting organization as well as IBM. One way to approach this challenge is to have easy and appropriate access to customers, the end-user, or potential consumer of your work. Our marketing affiliation ensures this, and we have numerous ways of reaching out and guaging customer response to our work and projects. Our feel for the real world also influences what we decide to work on, and management and staff spend considerable time on such decisions.

Many of the center's technical staff and project managers have served tours in systems engineering or marketing organizations. They understand the value of and need for on-going dialogue with customers during the life of a project. Because our work is viewed as experimental, we can be open with our visitors and are able to discuss and demonstrate as would most university departments. If and when our work evolves into proprietary or unannounced product areas, non-disclosure or contractual agreements become appropriate. We make every effort to minimize the need for these.

Organizationally, the centers have also enjoyed a dotted-line relationship with IBM's Chief Scientist, currently Dr. Lewis M. Branscomb, and his office. This arrangement is important to understand because in our view the centers are a corporate resource and are free to work across the corporation. We work and collaborate with other IBM laboratories in the interest of furthering the IBM cause as well as that of ACIS. Historically, corporate oversight of our affairs has come not only from the Chief Scientist but also the corporate Vice President of Marketing. Given the nature of the computing industry and IBM, this has been both beneficial and helpful to our work and in the conduct of our affairs.

THE TECHNICAL STAFF and the MANAGEMENT

I have already indicated that an important characteristic of a scientific center is its small size especially in comparison to other IBM laboratories. [Together, all three centers have approximately 100 professionals and managers]. Projects are also small and characteristically start with a single person (and not necessarily full-time) and may build to no more than three people or five at the largest. In sum, projects are intimate and focus entirely upon high quality people rather than organization and quantity of people assigned. Because projects usually go through a prolonged incubation period, they almost certainly start as the creative idea of one person.

The size of projects - they start small and stay that way - certainly colors the type of projects we undertake and places a premium on leveraging our small resource through collaboration, joint studies, and timely transfer to destinations outside our laboratory. This process is one of the trickiest and most time-consuming for our managers. Occasionally, questions of project viability and technical credibility arise when visitors view our staffing. While appropriate, such questions tend universally to reflect a lack of understanding of either center operation, the adtech process, or center accomplishments. (See the section ***"Case Studies from Cambridge"*** for examples).

The centers are typically located on or near a university campus rather than in the midst of a large IBM laboratory. While the objective is to have the centers involved in the professional and academic flavor of a university, each center also operates on a national basis. Because the centers are able to undertake joint work with university departments, it is not unusual to find the Palo Alto center undertaking a study with a department at MIT or Cambridge undertaking discussions with Stanford Linear Accelerator Center.

The university setting of each center also permits them to hire student employees to work on center projects. Our students work on the normal technical work and there are numerous instances of a student making a significant contribution to a significant project. In Cambridge students helped with the IBM Personal Computer prior to its announcement and a Cambridge student employee (now a professor at MIT) wrote the original text processor called SCRIPT which is still widely used on IBM's System/370. Finally, our university setting and student employees contribute significantly to the technical vitality and professional enhancement of the staff. In return we feel our students learn a lot about IBM close hand and many eventually go to work with IBM upon graduation.

The professional center staff has about 30 percent of its people with Ph.Ds but not necessarily in computer science. Where we consciously set out to hire at the doctoral level, we are typically looking for a computer science degree. At least 50 percent of the staff have an advanced degree, usually at the master's level.

Project managers for the centers come with considerable variety in their background. They may transfer in from another IBM laboratory or be promoted from the ranks of the technical staff. Not a few have IBM field experience, usually with systems engineering. Given our mode of operation and charter I personally believe this latter experience to be of considerable value. I look for and ask the project managers to be involved, player-coaches. Our centers and projects are too small to have the project management remote or aloof from the work. Our project managers have a relatively small span of control, averaging 5-7 staff members. Some may have as few as 2-3. We believe this is a reflection of the technical involvement we ask such management to assume.

Project managers and staff alike are asked to be flexible and feel comfortable crossing center project boundaries. This is essential since any staff member may at any one instant be involved in as many as three project activities which may typically be managed by several different project managers. Historically, this leverage has enabled us to respond rapidly to high-priority or crisis situations; seize opportunities of the moment that are important either to the center or IBM; offer the staff the continual ability to work at the important and/or leading edge; and feel there is an absolute minimum of jurisdictional and bureaucratic preoccupation. Such flexibility puts considerable pressure on everyone to manage the complex juggling act while seeing on-going work to completion. We are very conscious of the comparative ease in starting projects and the considerable anguish associated with completing/terminating work-in-process. This is truly a self-induced problem and one associated with an organization engaged in purely speculative work with a discretionary resource.

In terms of productivity, creativity, contribution, and track record of our people, I can find no correlation between tenure with IBM, degrees, background, or academic discipline. It is especially refreshing to find a relatively "junior" (i.e., recently-hired or -arrived) employee starting or leading one of our most promising or important projects.

UNDERTAKING the WORK

Perhaps the key managerial issue I face is that our work is all "local call." Being a discretionary resource, we are expected to initiate our own technical agenda within our laboratory and within the limits both of our resources and resourcefulness. In short, we follow our own lights.

In this sense the scientific centers are an organization accorded a level of freedom commonly associated only with individuals within IBM designated IBM Fellows. Like IBM Fellows, we both have tremendous freedom in selection of technical work combined with the frustrations associated with wondering whether our work is or will be appreciated. The trick is to capitalize upon the freedom and manage and contain the frustrations.

The result is a self-induced pressure to undertake relevant and, increasingly, near-in work. For our freedom hardly exempts us from the executive question -- implied or actually asked -- ***"what have you done for me lately?"*** So I tend to view my job as not unlike that of a university president who must contend with the matter of the endowment. For us, the endowment is one of contribution and achievement which allows me to minimize the past and focus or redirect the question to: ***"what are you going to do for us tomorrow?"*** In my view the thinking or concern of our management needs to be on how well we are running the race for the future. I try to get higher management to join me in concerns about what we are not yet working on -- but should be.

In the process, there is the risk that our past successes may obscure the real risk and effort needed for future success. Said another way, I work hard at trying to convey that regardless of whatever historic trove of golden eggs exists, the proverbial laboratory adtech goose does not deliver either on schedule or as a matter of guarantee. Past success can haunt you in this regard in my experience. Lack of past success, however, haunts you more.

As our laboratory management approaches the work, it is essential for them to possess a realistic understanding of IBM and its ways. In a dynamic company such as IBM, change is sometimes revolutionary but always on-going and a way of life. An accurate reading of the company is a precious and acquired skill. It is the single most important environmental skill and the one most relevant to achieving ultimate success.

But let us step back a bit and consider how things get started. First, we recognize that we are very much a function of who we are. In a real and positive sense, we are a captive of our in-house technical skills. I urge project management to simultaneously attend to matters of technical vitality and rotation of the staff for this reason, if no other.

Second, we work hard at avoiding so-called "forced marches" or projects imposed from on-high or from outside. This is not to say that patriotism and proper respect and deference to higher authority are lacking -- we just do not blend these elements into the technical agenda unless we already feel a strong positive inclination at work.

We feel most comfortable when our staff approaches project managers or myself with project proposals, ideas, or suggestions. In our view professionals have this responsibility and obligation.

Nonetheless, work starts in the laboratory in a variety of ways, and from the management viewpoint the process is a blend of negotiation, dialogue, and exotic minuet. I have already mentioned the preferred course -- the staff approaches local management. Based upon the track record of the staff member and other environmental, company, or industry factors, we may encourage the technologist to proceed immediately. ***"Keep us posted"*** is the last he/she is apt to hear from us for quite a while. We clearly like the idea and are betting on the individual. Others may get a flat ***"thanks but no thanks"*** if we are not enthused either with the proposal or prior track record of the proposer. Clearly, technical credibility and personal effectiveness counts in this process.

The middle ground is more likely to be a dialogue whereby management and staff reach an accommodation which results in a project more shaped by discussion and negotiation than in the first scenario. The result, however, is a viable project which both management and staff eagerly follow, perhaps, a little more closely than with scene

#1, above. This is the norm as well as the mean.

Less frequent are those situations where the professional (often senior) arrives at the door to announce: ***"The well is dry. What would you like me to work on?"*** We welcome capable exponents of such frankness with open arms for we all understand the on-going frustration of too many good ideas and too few hands at the bench. This individual is soon at work on an important but often anemically-staffed project. Frankness and objectivity regarding one's situation and state of affairs are encouraged in our laboratory and bring the bearer -- technologist or manager -- high marks. The proverbial Swiss Patent Office would also be a discombobulated turmoil if over-staffed with Einsteins.

Less desired are those situations where we must assign work to individuals. We insist upon it when a technologist is not productive over a long period or high priority or urgent matters come upon us and require prompt attention. In either case we take great pains to answer the ***"why?"*** involved with taking such a course of action. Then we do one other thing (and avoid doing another).

Usually the technologist on the assigned project gets caught up in it as it flourishes and takes on the flavor of the staff member as distinct from the original management aroma. We hear the individual talk of their project. Now the management must studiously refrain from any discussion regarding the origins of the work or ideas. In every respect it is now their "baby"! In sum, we feel work can occasionally be assigned to the staff, and as the work unfolds there is considerable management satisfaction flowing from this turn of events, and as we see the so-called "well" or "spring" flowing again.

Once a project is started, it may well fall into the low profile or so-called "under-the-table" category. While our primary management documentation consists of a quarterly project status report and an annual laboratory report, some work is so embryonic and exploratory we exempt it from these formal discussions and on-going scrutiny. Since everything we undertake is basically on speculation with our resource, we feel the responsibility for selecting the appropriate time to put the work formally on the table. We also feel a healthy laboratory has a split between these two realms at all times.

Nonetheless, a hallmark of any quality technical activity involves a significant start-up phase of homework. Literature searches and efforts to understand other peoples' work are crucial for avoiding reinvention of the proverbial wheel or for unearthing potential partners for collaboration. In this regard IBM possesses a remarkable asset which to our good fortune is an outgrowth of Cambridge work -- the VNET Network.

VNET forms one of the case studies appearing below, but for our purposes now we can say that it is a worldwide network linking IBM's technical community and individual technologists. It is both the world's largest network operating within one corporation as well as the electronic means for all IBM technologists to reach out to one another. Prototyped and developed at our laboratory, it has spread to some 3,000 computer systems throughout the company. VNET is a major factor in assuring us we are sharing our work as well as benefitting from that done elsewhere. In the fast-paced era which computer technologists live in, VNET has proven time and again to be an extremely powerful tool and our salvation in pursuing our work. To our amazement, we also find it is the way several hundred thousand IBM employees get their job done worldwide!

One last managerial item deserves attention before we turn to the actual pursuit of the work -- its various stages and fates. The item about to be discussed will also reappear in the next section but here it seems appropriate to comment on the absolute necessity for laboratory and project management to understand what is going on about them within their company, their industry, and the technical realm and field.

Were the managerial plate not already heaping, the necessity for this on-going alertness and managerial involvement beyond the project would itself be the occasion for the groaning helpings we ask our managers to herd back onto the project plate -- and then to contend with. The response to this stiff challenge, in my experience, distinguishes the various managerial performances I witness and evaluate, and has everything to do, using aquatic terms, with a manager's swimming versus treading water. As you might suspect by now, it has everything to do with successfully bringing our work to term and a happy fruition.

DOING the WORK -- Getting Down to Brass Tacks

When it comes to executing a project, the most common characteristic is that of **prototyping.** We actually build hardware or software and continually refine the technical approach rather than letting a study or strictly theoretical approach suffice. Prototyping, then, within IBM tends not only to be a hallmark of adtech, I believe it has an intrinsic and invaluable merit by itself.

Given the nature of technology discussions and decision-making within IBM, a prototype focuses the discussion, debate, or evaluation upon an existing, running, demonstrable entity which (1) can be seen and not just imagined; (2) measured and evaluated and not just discussed; and (3) altered and modified (if not improved) without feeling it has a life of its own or is incapable of change. A prototype is malleable and tractable -- anything else is usually less so. These are not new concepts or observations, but I feel they are a potent and effective approach. Let me give one example.

The idea for providing a System/370 and the Conversational Monitor System (CMS) on a desk-top originated in 1976 at Cambridge. This was long before serious personal computing and well before the IBM PC existed. As we pursued the technical implementation of our personal System/370, the words we dreaded the most were: ***"We believe you. We will take it as gospel that such a thing is possible. Why do you need to waste the time to build it?"*** But we did build it -- twice, and the final Motorola 68000-based prototype eventually became the basis for what IBM has announced as the PC/370. The prototype encapsulated our vision, highlighted problems and risks (performance) as well as advantages (M68000 capabilities and S/370 object compatibility). Having a running prototype helped us carry the day during the eight-year dialogue on what the industry now generically includes in the category of "super micro."

In general, adtech work and a prototyping approach can significantly cut product development lead time once a decision has been made to go forth. As an adtech facility we feel our job <u>is</u> to do our technical homework and provide IBM with alternatives and options. Using the example, above, it is interesting to note that on the same day, October 18, 1983, both IBM and Digit Equipment Corporation announced super micro implementations of their key architectures -- the IBM PCXT/370 and the DEC MicroVAX I.

From an environmental standpoint, the scientific centers do their work with the tremendous advantage of being remote from their reporting organization. ACIS and its director of scientific centers are headquartered in Milford Connecticut, and while there is not a little shuttling between this location and the three laboratories, the norm is blessed and relative relief from many distractions and the ability to do the work.

Yet there are visitations and seemingly interminable discussions of the work, but they are usually at the projects' initiation and a function of the openness with which we prefer to pursue the work. Management, nevertheless, remains sensitive as to when the level of discussion or travel are actually getting in the way of progress. Several other aspects of our openness deserve brief comment. First, everything is considered experimental and therefore open and subject to discussion. In this sense we parallel a research operation. Second, we are a "vanity publisher" and have a series of abstracted scientific center reports. The majority of these are in the public domain; a few are for "Internal Use Only" when proprietary issues are encompassed. We also encourage publication in outside journals.

Finally, we subject ourselves periodically to a peer review of particular projects. With the understanding of the technologists, I indicate that we are at some checkpoint in the work and may continue on, modify our direction, or terminate. We solicit frank opinions of the work and welcome collaboration. Such efforts pay huge dividends; force the project to crisp up its progress, status, and objectives; and usually attract collaborators from the invited auditors.

While peer review is a form of discipline and objectivity that the management deems appropriate, there is the obligation of project and laboratory management to continuously ask: ***"why are we still doing this project? What are we trying to accomplish? What is the end-game"?*** I recommend we ask technologists these explicit questions and sometimes the answers are quite revealing, unnerving, and disconcerting. ***"Because this is my life's work"*** is not an acceptable answer given the fact that we are not an academic or computer science department. But in the end I would far rather get that frank answer than none at all or for lack of never asking the question.

All of which is to say that the work has to be pursued with some level of vision. Something has to light the way. While serendipity and changing events certainly play an occasional role, a positive vision followed by concrete objectives and and a game-plan or strategy are absolute necessities. Like manners, they are conspicuous in their absence. Managers and technologists may, in fact, have differing visions or views of the work. This is normal and quite healthy, especially compared to the other extreme.

As we shall see in the next section, a heavy premium is placed upon two items throughout the life of the project. First, the absolute need to assess continuously the environmental factors affecting the fate and outcome of the work. Second, and closely related, is the early-on pursuit of the end-game -- technology transfer.

The END-GAME and SUCCESS

For us the end-game is a well managed conclusion of a project. We are in the business both of starting them but also obliged to write ***finis.*** Management skills show through more clearly in the ending rather than starting scenarios.

Two major routes constitute success for us and enable us to proclaim victory. In the most general sense, this is how we are measured.

First and most familiar is the quest for technology transfer. Since I have discussed how most of our work lives in a prototypal state, it follows that once feasibility has been proved and the prototype "flys", strong efforts are directed toward seeing the work leave our laboratory. It should go elsewhere in IBM in preparation either for internal use (such as with a tool or methodology) or for announcement as an IBM product.

Over the past ten years the vast majority of our work in Cambridge has seen the light as an IBM product. This can cause some confusion if not ambivalence when we consider that we are not ***per se*** a product development organization. We resolve this dichotomy through our conviction that top quality work deserves to be used by others outside our organization and that IBM's business policies increasingly require general availability through a product mechanism. Where an item transfers as an internal IBM tool, pressures may then build over time to let customers benefit as well. This two-staged approach has the disadvantage of taking considerable time but the very strong advantage of having passed the test of sustained and extensive use in-house. I find people like the idea of IBM taking its own medicine.

With this in mind, management makes every effort to see that our prototyping makes it as easy and possible for our transfer partner or "catcher" to pick up on the work. Where the following of standards or use of IBM hardware (current or future) makes sense, we go this route. Such has not always been the case with resulting impairment of the transfer, if not failure.

Another ingredient for success and an obligation for management is the early identification of and dialogue with a prospective transfer partner. Technologists show remarkable lack of interest in this and sometimes dismiss it either as **"selling"** or at the least an activity or skill more worthy of management than staff time.

Two points are in order here. First, the expected natural reaction of your prospect is to avoid you and show little or no interest in your overture. Adoption of Cambridge's project means, for them, the threat or reality of putting aside one of their on-going or planned activities. Thus the need for cultivation and selling <u>does</u> exist. Second, questions naturally arise in the center as to who engages in the selling/negotiation? Technologists usually try to assign this duty to managers, but we feel strongly that a partnership is in order -- a division of labor at the least. So we oblige the staff to participate and, if necessary, acquire new skills in the matter of presenting their work.

In the end, all we really ask is for the technologist to be able to convey the thrust of the project along with objectives and status. Professional pride inevitably leads to a fine showing by the people doing the work.

During the process of pursuing and cultivating a partnership, management and staff must be sensitive to the use of the word ***I*** versus ***we.*** Our prospective partner can easily be turned off initially or along the way if our inventors are insensitive and drum home the fact that I invented the ideas and project -- not you. Here the word "we" or "your/our" project yields rich human dividends. Our partners tend to have the perception that the scientific center gets the fun part or inspiration duty while the catcher or development laboratory is saddled with the perspiration phase of Edison's observation. There is more than a little truth to this view, and it behooves us to act accordingly and sensitively. It is not unusual to attract to our center one of our partners at the conclusion of a collaboration; this is a healthy and productive turn of events for us. It can also work in reverse which is both favorable and desired.

The unfolding of events in the prior discussion is the desired route to my self-styled end-game. While it may seen logical and straightforward and born of necessity, our transfer process is also protracted and not infrequently agonizing for all concerned. Only a strong conviction and major thirst for success by prevailing can sustain those involved over the many months or years of the cycle. In sum, our contemplative, remote, and obscure laboratory lives both in a world of reflective contemplation and upon a certain knife-edge of uncertainty. Sustaining morale and revisiting the larger view is not the least of our managerial challenges.

At some point the stiffest challenge arises when the desired route to transfer does not materialize -- or vanishes in mid-course. Either a partner eludes us or departs along the way. Then what? Now the question becomes a matter of whether we are willing to go it alone and assume the role of the developer, supporter, maintenance crew, and documenter. Were are thoroughly convinced these are skills we lack in any abundance or to any reasonable degree.

This leads us to the world of ***"it depends."*** Some of our work can never be released independently by ourselves, say in the case where we have made extensive enhancement and modification to an IBM operating system. Here the mainline developer and IBM "owner" of the operating system has to participate to assure quality, integrity, and to fulfill the IBM product commitment to many users.

In other cases we may, indeed, do the entire job where we feel it is in IBM or the work's best interest. For us, this is the stoniest road because it requires using the technical resource in a fashion that is hardly optimum and potentially counterproductive. Managerially, we take this route only with great reluctance; in few and far-between measures; and only after great internal soul searching and debate. Helping to fulfill an IBM commitment, rescuing an important corporate project; or to establish an IBM technology presence are the more obvious occasions to embark on this course of self-transfer.

Again, this places the premium on considering transfer early-on and while there is reasonable lead time. The proverbial last minute may be when the technical breakthroughs occur but it is the worst time to start considering partnerships, undertake negotiations, or expect transfer.

Some consideration is appropriate on motivation and reward of the technical staff. Universally, they express a conviction that the greatest reward is the opportunity to work on exciting projects and tackle non-trivial challenges. Everything follows from this point of departure. I explicitly commit to the staff that all levels of laboratory management will work to support their work in a quality fashion. We also will make transfer and obstacle-removal our priority business. That is the contract.

We also have in IBM a fine structure of professional and technical awards and our people are eligible. Because of our corporate mission, Corporate Outstanding Innovation or Outstanding Technical

Achievement Awards are appropriate for significant work. Where a sizable financial recognition accompanies these awards, the recipient and spouse may attend the Corporate Technical Recognition Event, usually hosted by the Chief Scientist. IBM support of our technical accomplishments has been extremely gratifying.

Yet the responsibility to initiate any recognition is appropriately on the plate of the project manager. I measure them on this item of stewardship. We also urge project managers to consider and utilize more informal avenues of recognition along the way during the various cycles of a project. There is a great disadvantage to storing up or reserving recognition for the ***grand finale*** -- especially with lengthy projects. How recognition is handled by the managers tells me a lot about their commitment to their people and work. It also lets me assess one more facet of line and second-level management effectiveness and creativity. I fully believe in recognizing significant project management effort with that of their staff. But such recognition is predicated upon the proper management job first being done on behalf of the staff.

In our daily operation there are inevitable tensions which require management attention. First, is the feeling by technologists that managers are inclined to discuss or promote the work prematurely. Frequently this is related to the feeling that the manager does not understand the work. The analog is management's concern over the technologist who never seems quite ready to talk about the work, discuss results or progress, or share his/her vision -- every "i" is being dotted and "t" crossed. Closet technologists and over-optimistic project managers always exist and both have a piece of experience to buttress their views. Mutual respect, open disagreement/discussion, and sensitivity carries the day. Neglect of this tension never does.

A second area for tension relates to the closing stages of a project. The innovation and creativity may be past along with the excitement; we are now going through transfer and wrap-up. The end is in sight. There are technologists who will seek to delegate this phase to others in their quest for the next, exciting activity. I refer to this outlook as:

not wanting to play the fourth-quarter of the game.

In my view it becomes a matter of professionalism if not keeping the faith, and it can require or consume considerable management attention and skill. Failure to navigate successfully through these troubled waters surely causes larger and ongoing problems down stream.

Finally, there can be tension if not paranoia among the staff on the matter of whether the managers are making commitments of the staff arbitrarily and without consultation. Just because this is not the case does not mean management need not be vigilant for signs of such concern. A track record of consistency, ethical dealing, and openness goes a long way toward addressing such concerns. Management also needs to work constantly to inform the staff on what is going on within the world (corporate and technical) beyond the laboratory. The broader staff knowledge is in regard to the larger picture and areas affecting them, the healthier and the better. Paranoia is frequently a handmaiden of a cloistered and constricted environment, and may be manifested by a ***"them versus us"*** outlook.

Perhaps the sternest test of all the managerial challenges covered in this section involves recognizing and seizing the opportunity for success. In essence it is the understanding that through a combination of circumstances which may be internal, external, planned or serendipitous, the vehicle for achieving project success has arrived. In unvarnished terms, this is known by us as a "pony." An alert and perceptive manager knows it for what it is and proceeds to ride it out of the gate. With such matters, he who hesitates **is** lost, and for want of its seizure the opportunity slips away. Whether an opportunity is of the project manager's own making or arrives at our door for a brief and transitory moment, the acid test is in what is made of it. Getting managers to think in these terms and incorporate it into their managerial kit is both a long-term and vital part of my job. The broader the project's vision and sharper their view of the work, the better prepared everyone is for such exciting and crucial moments.

In conclusion, there is another major way in which we are measured or by which we achieve success. This is the decision to terminate a project and move on. Acknowledging that we have failed and should either restart a project or end it is an important part of our managerial ethic. A premium is placed upon not calling a rose by another name. The focus is upon deciding what to do.

So we try to provide a climate where problems and short falls will be brought forward and fully reviewed. No time is wasted in the "woodshed" or in putting things under the rug or extracting them from under the rug. Some of our healthier work has risen phoenix-like from the ashes of false starts.

I believe this is one of the hallmarks and blessings of an adtech activity, and one of the advantages of working in prototype media. Short falls and failures are one indication to us and our management that we are reaching out and stretching -- not relying on comfortable skills and past triumphs. Knowing when a project or person is floundering or when to call it a day is as important as the other measures of and avenues to success we have already observed.

CASE STUDIES from CAMBRIDGE

I. Distributed System/370

In 1976 Cambridge initiated project work to investigate the distribution of System/370 architecture and applications. The motivating force was the conviction that hardware technology trends would enable System/370 to be packaged in increasingly smaller size, thereby enabling this popular computer family and its application software to reside outside the traditional computer room. Clearly this is what was already happening with the minicomputer phenomenon -- but without remote support and not for System/370.

The objective was to explore and prototype a System/370 which could physically operate in a laboratory, business, or office environment (to pick some examples). The challenge was to determine if such systems could be supported and even operated by a technical staff remote from the computer where organizations desired this approach to distribution. Two prototypes emerged pursuing these goals from distinctly different directions.

- ***Remote Operations of System/370:***
  The problem here involved work on software and hardware vital to operating and supporting a small System/370 in an workplace environment -- with only end-users present. The challenge was to provide neces-

sary software and computer operations support remotely to a System/370 which, for all purposes, was unattended. The model was very much that of telephone switching equipment -- the end-user would access the computer application via a terminal but remain unconcerned and uninvolved in where the small System/370 was located (down the hall, in a closet, in the corner?) and how it was supported. This experiment, in effect, was working on the economics of distributed computing by allowing one set of technical specialists within an organization to manage, control, and operate many remote systems from which remote users were being serviced for their applications.

Initially, experimental work was done on the smallest System/370s which, at this time, did not lend themselves to distribution or remote operation. Work progressed using a microprocessor-based special computer attached to the small System/370. It would remove the necessity for the traditional console operator at such a system. Console operators read messages from the operating system and its end-users, type responses or issue commands, and push buttons on the system. Our experimental micro would be on a communications link to a central site and would instruct the distributed System/370 in lieu of a human, local operator. Our conviction grew that such a micro should be an integral part of future System/370s.

Software work involved enhancements to the VM/370 operating system to recognize that messages and commands requiring human operator action had to be sent on a network from the distributed system to the actual human operator at a central site.

Work progressed steadily throughout 1976-77 at which time we reached a major fork in the road. Opportunity or the so-called "pony" was at hand as we learned that IBM product laboratories in Germany and Endicott, New York were undertaking development of minicomputer-size versions of System/370. They were eventually announced in January 1979 as the 4300-series. Having learned of this activity, Cambridge management now had to decide whether to take time away from the adtech work to form a partnership with the 4300 developmental engineers to ensure these machines would incorporate our experimental work. The issue for us was to achieve the transfer of the Cambridge work into the 4300.

In the end, we did just this and our technology for distribution was adopted by the various 4300 development laboratories for inclusion in the 4331, 4341, 4381, and 4361 processors. Our running prototype and experimental head start in 1977 helped considerably in the debates that took place over the 4300. The homework had been done in anticipation of the situation that unfolded in a major IBM product line.

By 1983 the last of our remote operations work had been announced. It provides the 4361 the unique capability of powering itself on or off at a designated time and/or day. Likewise it can do the same in response to a telephone call from anywhere in the world or in response to restoration of electrical power after an outage. These last capabilities bring to mind another facet of how we pursue our work -- joint studies.

We are able to undertake studies with partners outside IBM, and in the case cited above we undertook a study with the University of Maine to test our work in a "real world" situation. Here the objective was to assess the remote 4300 hardware and software in a live environment with a small 4331 at the University of Southern Maine (in Portland) operated and supported from the main computer center at Orono (140 miles away). For more than a year the Portland machine ran unattended, providing timesharing services to faculty and students. Yet their system was operated, controlled, and maintained from the Orono campus. This test provided Cambridge on-going assessment of our experimental work and valuable development feedback to the 4300 product engineers.

- ***Desk-Top System/370 -- PC/370:***

For our project at Cambridge, the local end of distributing System/370 involved placing a single-user version on the individual's desk. This was in 1976 long before the implications of the micro revolution were clear or today's personal computer was born.

Because of its interactive nature and wide acceptance on System/370, Cambridge focused upon placing the Conversational Monitor System (CMS) portion of the VM/370 system into a personal system. Over a period of several years a prototype was built using a special version of the IBM 5110 portable computer provided by the Rochester, Minnesota laboratory. By the time we had a demonstrable system it became clear to us that performance issues had to be solved before serious attention could be gained in IBM. Also, the unfolding microprocessor revolution indicated to us the folly of continuing to use a proprietary IBM micro such as existed in the 5110. We needed to avail ourselves of outside technology such as that rapidly emerging from Intel and Motorola in their standard microprocessor sets. IBM was starting to do the same.

In parallel with our technical work we aggressively lobbied within IBM for the capability we were then prototyping -- a CMS application delivery system available to individual users on their desk. Having the prototype was immensely helpful as we not only talked but showed the concept in operation. In 1981-82 when the Motorola M68000 microprocessor became available to us, we redid our entire prototype to place us in a state-of-the-art technology environment. This greatly increased the credibility of the work while enormously increasing the workload on the project as they trekked back up their the prototypical mountain.

Again, opportunity presented itself. In 1982 a consortium from the Yorktown Research Division and the General Technology Division laboratory in Endicott reviewed our work and decided to adopt our prototype as the basis for a serious IBM product based upon the newly-arrived IBM Personal Computer. Specifically, the PC/370 was to involve use of M68000 coprocessors and the CMS operating system on the forth-coming PC/XT and PC/AT to provide System/370/CMS application compatibility as an option for hard-disk models of the PC. Cambridge transferred its work to the consortium which embarked upon a product effort, first announced in October 1983 as PCXT/370. In Cambridge we saw our role as one of invention and innovation with the additional challenge of helping IBM decide to include the System/370 option for the IBM Personal Computer. Thus far in 1985, IBM has also announced PC options for its Series/1 and System/36 while elsewhere Digital Equipment Corporation has done likewise with its MicroVax I & II.

II. ASYNCHRONOUS (ASCII) FULL-DUPLEX COMMUNICATIONS

For many years IBM has offered System/370 systems with a terminal family known as the 3270 and which utilizes an IBM protocol. The 3270 family of terminals is characterized by its high-quality and relatively high expense (i.e., over]1,000). Because of its university and research contacts, Cambridge was concerned that IBM did not provide either a terminal or operating system support for the Asynchronous ASCII Full-Duplex protocols. ASCII is very common on campus and is especially attractive because of the relatively inexpensive terminals (i.e., under]1,000) implementing it. Cambridge efforts focused upon two distinct arenas for enhancing IBM's posture in regard to ASCII communications.

- ***ASCII Terminal and Attachment to System/370:*** In 1978 Cambridge became aware that IBM's FUJISAWA, Japan laboratory was undertaking development of IBM's first ASCII display. Concerned that this terminal embody the extensive knowledge found upon college campuses where ASCII was already well-understood, the center facilitated consulting arrangements between a number of university computing centers and developmental engineering in FUJISAWA.

At the same time, Cambridge focused upon the problem of how IBM's ASCII terminal (announced as the IBM 3101) would attach to System/370 mainframes. The objective was to allowed inexpensive ASCII terminals to access IBM mainframes without changing the application programs running on System/370. Through the efforts of a Cambridge project manager, a partnership was formed between IBM development laboratories and the Yale University Computer Center to utilize IBM's Series/1 minicomputer as a protocol converter. The result was a program offering from Yale and marketed by IBM for full-duplex ASCII support for System/370. IBM also made ASCII enhancements to its VM/370 operating system after much Cambridge lobbying. The work from Yale has become one of IBM's most wide-accepted products in the university and scientific arena, and IBM has recently announced several ASCII protocol converters to improve further the economics of such terminal attachment to System/370.

The center's ability to perceive the real world situation and pursue a solution with a variety of partners -- inside and beyond IBM -- serves as an example of the role and mission and usefulness of the scientific centers. In effect, we feel the scientific center played a pivotal role in helping IBM change, then enhance its product offerings in an area the scientific centers offer long-time expertise.

- ***ASCII Communications for the IBM PC:*** During the debates over the PC, Cambridge played an active consulting role. One of the areas of greatest involvement related to communications for the PC.

Cambridge felt strongly that (1) the PC should communicate both to other PCs and larger computers, and (2) the PC should be able to communicate to environments other than those of IBM's. In our mind this placed ASCII communications at the head of the list for modes of communication for the PC.

After no little discussion in the spring of 1981, agreement was reached on announcing the PC with an ASCII communications card. Now the question became one as to whether there would be IBM ASCII software support on the PC. Given a narrowing window for announcement, the answer seemed to be ***"no."*** The ball also landed in Cambridge's court.

Given our long campaign for ASCII within IBM and strongly-voiced convictions on PC communications, the challenge was whether we would undertake the development task of providing a PC package. After considerable soul-searching in the laboratory, we made a commitment to the Boca Raton development group to provide an ASCII communications package. It was announced in August 1981 as "Asynchronous Communications Support" and was available in December 1981. This package has undergone two revisions and been joined by a full-duplex ASCII package "IBM 3101 Emulator" for the PC.

These packages are outstanding examples to us of a scientific center helping to meet a significant IBM need in record time. They are also unique because of the extremely low "bug" rate, advanced coding techniques employed, and significant customer acceptance. Together, these Cambridge packages have reached over 100,000 users. Two technical staff members and one project manager were the Cambridge resource, and this hearty band has the satisfaction of making a major contribution in true scientific center tradition.

III. VNET NETWORKING

In 1975 when I arrived in Cambridge there was a small one-man project involving software to network the three scientific centers and Research Division laboratories in Yorktown and San Jose.

Our technologist secured the part-time services of a colleague in Poughkeepsie and before long our five-node store-and-forward network started spreading to other IBM laboratories. To make a long story short, today this networking facility which we call VNET links over 3,000 computer systems within IBM and was transferred into the VM/370 operating in 1977. It is now used on over 5,000 IBM systems. Totally justified as an internal IBM tool, VNET has also become a major IBM product.

VNET is also the technological underpinning for a fast-growing university network outside IBM called BITNET. BITNET now has over 500 nodes connected within the United States and has connections to two other networks which are also based upon VNET. NorthNet connects a growing number of Canadian universities and European Academic and Research Network (EARN) connects many universities in Europe and the Middle East.

Cambridge has the satisfaction of seeing an effort of some 1.3 equivalent staff unfold within IBM and without in this bewilderingly rapid fashion. This, then, is probably our most pervasive and widely-accepted adtech effort judged by the number of people and institutions affected. I say this even given the fact that IBM's VM/370 operating system

originated as an adtech project in the 1960s
in the Cambridge Scientific Center.

SUMMARY & CONCLUSION

The scientific centers are a very small resource within IBM's United States technical community. They engage in a very specialized activity known as advanced technology which is best characterized by the building and evaluation of work through prototyping.

Work proceeds on a bottoms up basis which places a premium upon project selection and on-going evaluation of work underway. Because the centers are a discretionary resource, are allowed to work across the entire corporation with corporate oversight, and may also undertake work which will occasionally be counter to IBM's main strategies, emphasis is placed upon openness and extensive communication of technical activity. As illustrated in the previous section, results can be significant to the corporation and represent direct response to existing or emerging problems. Anticipation of problems and recognizing avenues and moments of opportunity is crucial to the centers. Breaking new ground is just as important as addressing near-in problems.

A combination of technical staff excellence and involved project management are the basic building blocks for projects. Project managers are expected to focus upon the business of transferring or terminating work and urged to enlist the involvement of their technologists in this process. Attention to developments outside the center is crucial for all facets of the laboratory's work. This is an on-going and demanding challenge -- not an explicit phase of the work.

All parties in the center have the potential satisfaction of working on real, significant problems; making a major contribution to IBM's technology agenda; and the not infrequent occasion to see their work flourish in practical use in the real world. These satisfactions are the positive counter weight to the inevitable frustrations associated with long-term, demanding technical work along with the on-going concern as to whether center activities are properly appreciated or understood. Such is our life within an industrial organization.

The centers have lived their existence for nearly twenty-five years without substantial change in their charter, mission, organization, or size. Occasional fine-tuning has occurred and specialized centers have closed. The current three centers seem to represent steady state over the past ten years and focus largely upon system and programming technologies or science and engineering applications.

My feeling is that major changes (such as a dramatic increase in headcount or top-down project direction or more emphasis upon product development) would not improve the centers nor enhance their usefulness to IBM over the long haul. A delicate and poorly-understood balance exists in the centers' operations, and I have a very healthy respect for the results that been achieved over a very long time.

The basic anomoly, then, is this somewhat disconcerting balance. The centers exist somewhat at odds within a corporation known for its dedication to crisp organization, constructive change, and extensive measurements. Yet they have produced a distinguished endowment of well-received and widely-used technologies for an incredibly modest investment.

The Palo Alto Scientific Center's Contributions to Languages and Programming Technology

Harwood G. Kolsky

The Palo Alto Scientific Center has for twenty years provided an atmosphere in which highly individualistic, talented and often eccentric software and hardware experimenters can flourish. The emphasis has been primarily on practical prototypes, some of which have become very successful IBM products. Others must be classified as "lost opportunities." Many awards have been given for the projects mentioned here, including three Outstanding Contribution awards. The following are brief summaries of the projects:

o Microprogramming

Microprogramming as originally used in System/360 was a convenience to enable engineers to emulate the S/360 architecture on several different internal mini-computers. Up to that time the "programs" tended to be short since only a few microinstructions were needed to emulate most S/360 instructions and each emulated instruction was a separate program.

In 1968 A. Hassitt and H. S. Smith set about to demonstrate that a high level language (APL) could be emulated directly. The machine chosen was the S/360 Mod 25 (Ref. 2 and Ref. 12). J. W. Lageschulte and L. E. Lyon also worked on the Mod 25 project (Ref. 4). The work was generalized (Ref. 3) and later turned into high-performance products for the S/370 Mod 145 by Hassitt and Lyon (Ref. 7), and for the S/370 Mod 135 by R. G. Scarborough, H. G. Kolsky and N. S. Gussin (Ref. 20). The projects used different design approaches because of the different underlying minicomputers in the machines. Execution performance gains of 10 to 1 over the existing software APL systems were common. Because of the way IBM viewed microcode, the microcoded assists were considered as "machine features" and not as program products.

Given its early availability, the APL assist served to demonstrate what an effective microcode assist can achieve, and it stimulated subsequent assist implementations.

The Palo Alto microprogramming efforts would have been very difficult without the use of two special compilers written by J. R. Walters Jr. and D. L. McNabb (Ref. 5 and Ref. 8). These compilers permitted the program developer to write microcode in a high-level PL/I-like language. The compiler generated assembly language in the correct formats for input to the microcode assemblers for the respective S/370 models.

o APL Interpreter Development

An APL system which used the APL emulator began running reliably in the summer of 1972 and was distributed for use within IBM. This system could run only on the S/370 Mod 145 since it used the 145 microcode. A decision was taken at Palo Alto to produce a compatible software system. The resulting system, subsequently named APL/CMS, was completed early in 1974 and announced for release to IBM customers in May 1974 (Ref. 17). The work was done by R. J. Creasy, M. Beniston, A. Hassitt and L. E. Lyon. APL/CMS formed the basis for the VSAPL Program Product. Changes were made to the APL/CMS supervisor but the VSAPL interpreter used all of the code from the APL/CMS interpreter.

APL/CMS could run on any IBM System/370 running VM. When running on the 135 or 145 it would utilize the microcode emulators mentioned earlier. Another project to provide a microcode assist for the 4331 was also successfully completed but dropped because the floating point hardware on that machine was too slow to benefit from it.

The development of APL/CMS lead to many new techniques for implementing APL, in particular the syntax scan (Ref. 18) and the efficient evaluation of array subscripts (Ref. 9). The software version was typically a factor of two faster than existing APL systems. With the microcode, it was another factor of five faster in a significant range of cases. This work was also unique in that it resulted in one of the few patents ever granted to IBM on program language technology (Ref. 6).

o VSAPL Interpreter Performance Enhancement

In 1979 R. G. Scarborough began a project to develop an experimental version of the VSAPL interpreter designed for high performance. The intent of the project (named VQAPL) was to gain in software alone a performance equal to that one might obtain with special microcode. The performance gains came from both a general cleanup of data handling and register usage, and from new methods of performing some of the key tasks, such as space management, scanning, etc.

The results of this work were impressive: The new interpreter required only one-third to one-half as many machine instructions to execute an APL primitive function compared to the VSAPL interpreter. However, for non-technical reasons this work was never incorporated into VSAPL.

o APLGOL

APLGOL is an interesting offshoot of APL. Originally developed by R. A. Kelley and J. R. Walters in 1972 (Ref. 13 and Ref. 14), it replaces the primitive go-to (right arrow) of APL with an interstatement control structure such as if-then-else or do-case. The APLGOL compiler is unique in that a common internal form is kept for the APL and APLGOL text. This permits consistent modifications to be made at either level (Ref. 16). It continues to be used for exploratory work.

o Decision Table Compiler

In the early 1960's M. Montalbano did some of the first work published in the field of decision tables. Decision tables are a rigorous way of specifying conditions in the description of procedures (Ref. 1). He expanded on this work during his many years at PASC. In 1974 his book entitled "Decision Tables" was published by SRA (Ref. 15).

In 1975 H. J. Myers developed an optimizing decision table compiler, DTABL. It took input in the form of decision tables and generated the controls to produce complete application programs. The optimization techniques are described in Ref. 11. DTABL is written in APL but prepares programs for use in PL/I, COBOL, or APL. These can be sent to the appropriate compilers using the shared-variable interface. The program was an brought out as an Installed User Program in 1976 (Ref. 21).

Recently S. M. Roberts, H. J. Myers and H. F. Smith have reexamined some of the theoretical conditions of decision tables and have extended their usefulness (Ref. 37).

o General Cross Assembler Generator

In 1976 H. J. Myers discovered and implemented a technique by which an assembler written in APL could perform at speeds comparable to those written directly in assembly language (Ref. 19). This technique was used to develop an Installed User Program which has been particularly useful in producing cross-assemblers for mini- or micro-processors (Ref. 22). There have been several examples in which a cross-assembler for a mini-computer has been written in less than a day using the product.

o 5100 Event Driven Executive (EDX) and Basic Compiler Prototype

The IBM 5100 computer was based on a PASC-Los Gatos hardware prototype called "SCAMP" originally developed by P. J. Friedl, P. A. Smith, and J. D. George in 1973. SCAMP was IBM's first portable computer and was the direct ancestor of the PC (Ref. 36).

In 1976 an investigation was undertaken by C. R. Hollander and H. F. Smith to determine whether the performance of the 5100 with Basic could be increased significantly by implementing the System/7 EDX supervisor in 5100 microcode, and compiling Basic source programs to EDX target code. The work was technically a success but was not made into a product (Ref. 27).

In 1978 P. A. Smith and N. S. Gussin wrote a user's guide for their 5100 Even Driven Executive system (Ref. 26). The 5100 EDX was a programming tool patterned after the S/7 and Series/1 EDX's which were developed in San Jose Research. The program utilized a set of macro commands to facilitate the generation of multiple application programs which could operate simultaneously. It had a sensor-based front end, and provided commands for data acquisition, reduction, and analysis, as well as for report generation. It provided a method for the rapid implementation of new applications.

o Array Language Emulator

Trenchard More of the Cambridge Scientific Center has been developing his Array Theory since 1968. By 1976, More had developed an impressive theory and an axiomatic basis for defining operations on arrays. At this point, A. Hassitt and L. E. Lyon began work on implementing More's theory with a view to developing a useful computer language. Since the final form of the language was not yet known, they implemented an "array theory machine" in software (Ref. 24 and Ref. 25). The machine could carry out More's operations on arrays but did not impose any language aspects. A shared variable interface between the machine and APL was built so that an experimental language could be developed (Ref. 28).

The work proved significant in a number of respects. It revealed some gaps in the axiomatization of the operations which More was able to fill. It enabled More to experiment with the array theory operations and develop the theory in a direction which would be most suitable for practical application. It developed new methods in implementing the complex data structures which can arise in array theory. It enabled Professor M. Jenkins of Queen's University to develop an array theory language, using the shared variable interface to the array theory machine. The experience which Jenkins gained in that project enabled him to design a language called Nial which provides a powerful language which is well matched to the power of More's array theory. Jenkins, with major grants from IBM, and More have continued to develop Nial.

o Efficient Use and Conversion of FORTRAN Programs

Although more exotic languages occupy the center of attention, PASC has always realized that the "bread-and-butter" language of engineering and scientific computing continues to be FORTRAN. Much of the value of Center's work has come from its support of the field in this area. In 1972 R. N. Kortzeborn wrote a popular report on how to make efficient use of large computers from FORTRAN (Ref. 10). In 1979 A. A. Dubrulle wrote an excellent tutorial on the design of matrix algorithms for FORTRAN and virtual storage (Ref. 38).

More recently R. A. Blaine has provided an overview of FORTRAN on IBM systems, discussing the real problems of conversion and portability (Ref. 34). There have been more requests for this report than any other PASC language report.

o FORTRAN H Extended Optimization Enhancement (FORTRAN Q)

The original versions of the IBM FORTRAN H compiler were written between 1963 and 1967 to support the IBM System/360. FORTRAN H was considered the best compiler in the industry for many years. In 1976 a small study was started at PASC to determine why certain inner loops in FORTRAN programs were not optimized consistently by FORTRAN H. As a result of these findings, Randy Scarborough initiated a project to produce a new optimizer.

The strategy in building the new compiler was to incorporate the existing routines which worked well, rewrite those which did not, and write new routines to perform additional optimizations. Later improvements were made to the FORTRAN library as well (Ref. 29). The results were made available in September 1978 as an Installed User Program (Ref. 23).

Randy Scarborough's now famous FORTRAN Q compiler was a key step in reviving IBM's image in the Scientific marketplace and continues to receive excellent customer acceptance. FORTRAN Q Version 1 Mod 4 was released in May 1983. This release includes support for the FORTRAN Interactive Debug.

o FORTRAN Interactive Debug (IAD)

In response to pleas from the field in 1981, R. G. Scarborough began a project to provide support for IBM's current FORTRAN Interactive Debug (IAD) program product in the then-new VS FORTRAN compiler.

A modified FORTRAN Q compiler was built which interfaced with the IAD at the FORTRAN '66 level. It was then modified further to extend the interface for the FORTRAN '77 level data types. The modified compiler thus generated test input equivalent to the real input expected from the VS FORTRAN compiler. Next the IAD was modified to accept the FORTRAN '77 language extensions. Finally, the changes made to FORTRAN compiler and the IAD were transferred to GPD and formally installed by them into the program development process. The resulting product was made available to customers in January 1982.

o FORTRAN Utilities for VM/370

In 1981 P. M. Hirsch prepared a very useful set of support utility routines to assist FORTRAN programmers. It was developed to meet or exceed the specifications required by the electric power industry for their engineering applications (Ref. 31). The program subroutines were designed to be called from either FORTRAN '66 or '77 language standard programs. The package supports many data transfer operations, data set operations, etc. directly from FORTRAN programs.

o ADRS/II Development

John McCleary had the original idea for an Administrative Data Report System (ADRS). He implemented the original version in APL while he was a DPHQ staff member in White Plains. Greco, an SE in a New York branch office, extended it to handle file I/O. As an IUP, ADRS was successful to a degree, but suffered from a performance problem.

In 1978 McCleary contacted H. Joseph Myers of the PASC IBM Fellow Department, who agreed to design and implement a calculation-statement compiler for ADRS. Myers' contribution also included revision and extension of the original calculation language of ADRS, as well considerable consulting on matters of performance design of several of the ADRS/II components. The resulting package was issued in 1979 as a new IUP (No. 5796-PLN) with about ten-fold performance improvement. It completely replaced the original program.

Because it is a perpetual IUP, Greco and Myers have continued to provide consulting services for problems encountered by the maintenance teams. In 1982 Greco enhanced the product with a color graphics feature.

o PASCAL Compiler and Interpreter for Series/1

With the goal of making Pascal available for Series/1 users, L. M. Breed and H. F. Smith in 1981 (Ref. 32) designed and implemented an execution environment including storage organization, procedure linkage, and system services. Emphasis was on space and time efficiency. Code optimization was at the peephole level, i.e. no global or inter-statement optimization was done. The compiler was written in Pascal. It ran on a System/370 under VM and produced Series/1 object code decks in two passes. This code was then down-loaded to a Series/1 via a communications line and catalogued on disk, where it could be linked with run-time support code and executed.

During the same period, A. Hassitt, N. R. Mendelsohn, P. Sih, and P. A. Smith were developing a Series/1 Pascal interpreter (Ref. 30). The interpreter and the compiler used the same run-time storage organization and procedure linkage, along with many common system components. L. E. Lyon also worked on the run-time I/O support for the S/1.

o Pascal/VS for the Displaywriter

The UCSD (University of California, San Diego) P system is available for the Displaywriter, and it provides a Pascal compiler and interpreter. However, the resulting code must be run with the P system and cannot be used with the standard Textpack system of the Displaywriter. In 1981, A. Hassitt and L. E. Lyon developed software for running Pascal/VS programs under the Textpack system. The software consists of (a) a new third phase for the Pascal/VS compiler, (b) re-compilation of programs from the Pascal/VS run-time library, and (c) new routines, some written in Intel 8086 assembler code. In 1982, a debug feature, some performance improvements, and the changes necessary to run the software on the latest version of Textpack were made (Ref. 33).

o Pascal/VS for the Personal Computer

The microcomputer on the IBM Personal Computer is an Intel 8088. Since the Intel 8088 has the same instruction set as the Intel 8086, some of the above Displaywriter code could be used unchanged on the Personal Computer. In 1982 the complete new system was finished and called Pascal/PC (Ref. 35).

Pascal/PC has the advantages that it uses the Pascal/VS language and many Pascal/VS programs will run with very minor or, in some cases, no changes. It supports 32-bit integers whereas Microsoft Pascal supports only 16-bit integers. UCSD Pascal supports long integers but with several important restrictions. Pascal/PC supports strings of up to 32767 characters whereas the other Pascal compilers are limited to 256 characters. The performance of Pascal/PC code is reasonable when measured against other interpreters but is 10-20 times slower than compiled code.

o Vectorizing FORTRAN Compiler (TVFtran)

Starting in 1981, a joint study between PASC and Prof. Ken Kennedy of Rice University investigated the development of automatic vectorizing programs which can produce correct vectorized output from source programs written in ordinary FORTRAN. Randy Scarborough incorporated the resulting algorithms into an experimental vectorizing compiler.

The first version of TVFtran became operational in April 1983. This interim version successfully compiles vector object code from unmodified scalar FORTRAN programs. It incorporates in a working FORTRAN Q compiler Randy's earlier work on vector register allocation and new versions of the Rice University vectorization algorithms. It is a solid base for expanding the capabilities of the vectorizing compiler in the future.

The TVFtran compiler presently converts do-loops into vector object programs only when the do-loops are entirely vectorizable. The main work in process has been to remove the restriction that the entire loop must be vectorizable before any vector object program will be generated.

Much progress has been made in developing and programming an algorithm which will: identify substantially more opportunities for using the vector hardware on a nest of loops; compare the costs of executing statements along the various possible dimensions; determine when statements may be merged into common controlling loops; determine the minimum arrangement of controlling loops to execute the merged vector program; etc. The remaining important task to be done is the vectorization of loops containing if-statements.

o Hybrid Vectorizing Compiler (VS FORTRAN and TVFtran)

The most practical method for incorporating the improved vectorizing optimizer in an IBM Program Product is to delete the entire optimization phase of VS FORTRAN and to install in its place the optimizer, written in FORTRAN, intact from TVFtran. There are some changes in the data structures, but the body of the optimizer routines is unaffected. In December 1983 the optimizer from FORTRAN Q was installed in the VS FORTRAN compiler. This hybrid produces identical object decks as the current PID-level VS FORTRAN Release 3. Assuming everything goes well, the final step will be to incorporate the TVFtran vector changes into the hybrid compiler.

We do not feel that we need to argue whether or not the resulting compiler, part in FORTRAN and part in PL/S, is better than one written entirely in PL/S. We are interested in a practical solution to problems and are willing to leave the "metaphysics" for later.

o Extended Exponent Extended Precision Arithmetic (XEXP)

Ever since System/360 was introduced twenty years ago the "limited exponent range" of the IBM floating-point number format has been troublesome in special marketing situations. In 1983 A. Karp, A. Hassitt, and R. G. Scarborough found a deceptively simple solution. The FORTRAN Q compiler was modified so that, under control of a user option, every arithmetic operation involving extended-precision operands would be changed from a hardware instruction into a subroutine call with special fast linkage. Subroutines were written to allow extended precision exponents to range from 10**-9865 to 10**9861.

This work has been completed and two important features emerge. First, the expanded range is available directly from the FORTRAN source language by specifying REAL*16 or COMPLEX*32 and requesting compilation with the new XEXP (for extended-exponent extended-precision) option. Second, the compiler itself does not know anything about the arithmetic being implemented by the subroutines. The subroutines could chose to implement interval arithmetic or IEEE floating-point arithmetic instead of S/370 XEXP arithmetic. This is a way to support arbitrary arithmetics from FORTRAN source without changing the compiler.

REFERENCES

1. M. Montalbano, "Tables, flowcharts and program Logic," IBM Systems Journal, vol 1, pp 51-63, (Sep 1962).

2. A. Hassitt, "An Assembler for IBM System/360 Mod 25 Microcode", PASC Tech. Report 320-3268 (Sept 1969)

3. A. Hassitt, J. W. Lageschulte & H. F. Smith, "A Procedure-Oriented Machine Language", PASC Tech. Report 320-3271 (Jan 1970)

4. A. Hassitt, J. W. Lageschulte, & L. E. Lyon, "A Microprogrammed Implementation of an APL Machine," PASC Tech. Report ZZ20-6402 (Jul 1970).

5. D. L. McNabb & J. R. Walters, Jr, "MPL/145 A Language and Compiler for System/370 Mod 145 Microprogramming", PASC Tech. Rpt. ZZ20-6410 (May 1971)

6. A. Hassitt, "Method and Apparatus for Operating Computer Apparatus in Response to Program Statements", Patent Appl. 129,739 (Mar 31, 1971)

7. A. Hassitt & L. E. Lyon, "An APL Emulator", PASC Tech. Report ZZ20-6417 (Jul 1972) also a revised version, "The APL Assist", ZZ20-6428 (Feb 1975). Externally published as:
A. Hassitt & L. E. Lyon, "An APL Emulator on System/370", IBM Systems Journal, Vol. 15, no. 4., pp. 358-378 (1976)

8. D. L. McNabb & J. R. Walters, Jr., "MPL/135", PASC Tech. Report ZZ20-6415 (Mar 1972)

9. A. Hassitt & L. E. Lyon, "Efficient Evaluation of Array Subscripts of Arrays", IBM J. of Res. & Develop., vol 16, pp 45-57 (Jan 1972)

10. R. N. Kortzeborn, "Efficient Use of Large Scale Computers from FORTRAN," PASC Tech. Report 320-3286 (May 1972).

11. H. J. Myers, "Compiling Optimized Code from Decision Tables", IBM J. of Res. and Develop., Vol. 16, pp. 489-503 (Sept 1972)

12. A. Hassitt, J. W. Lageschulte & L. E. Lyon, "Implementation of a High Level Language Machine", Comm. ACM, Vol. 16, pp 199-212 (Apr 1973)

13. R. A. Kelley, "APLGOL, an Experimental Structured Programming Language", IBM Journal of Res. & Develop., Vol. 17 (Jan 1973)

14. R. A. Kelley & J. R. Walters, Jr., "APLGOL-2, A Structured Programming Language for APL", PASC Tech. Report 320-3318 (Aug 1973), also publ. in Proc. of APL-VI Conference, (May 1973)

15. M. Montalbano, "Decision Tables," book publ. by SRA Inc. (1974).

16. H. G. Kolsky, "APLGOL, A Structured Programming Language for APL", Programming Methodology, ed. by G. Goos & J. Hartmanis, p. 161-191, Springer-Verlag (1974)

17. IBM Program 5799-ALK, IBM PRPQ MF2608 and IBM RPQ 500256, "APL/CMS and APL Assist," (1974).

18. A. Hassitt & L. E. Lyon, "APL Syntax and APL Execution", PASC Tech. Report ZZ20-6432 (Dec 1976)

19. H. J. Myers, "A Fast Assembly Technique using APL", PASC Tech. Report ZZ20-6431 (Jun 1976) also published in IBM Tech. Disclosure Bull., Vol. 20, No. 1 (Jun 1977)

20. R. G. Scarborough, H. G. Kolsky, & N. S. Gussin, "Microcoding APL: The System/370 Mod 135 APL Assist", PASC Tech. Report ZZ20-6433 (Jan 1977)

21. IBM Installed User Program 5796 PJB, "APL Decision Table Compiler", (1977)

22. IBM Installed User Program 5796 PKD, "General Cross Assembler Generator", (1978)

23. IBM Installed User Program 5796 PKR, "FORTRAN H Extended Optimization Enhancement," (1978).

24. A. Hassitt and L. E. Lyon, "Array Theory/370," PASC Tech. Report ZZ20-6435, (Jan 1978).

25. A. Hassitt and L. E. Lyon, "A PL/I Program for Experimenting with Array Theory/370," PASC Tech. Report ZZ20-6436, (Jan 1978).

26. P. A. Smith and N. S. Gussin, "5100 Event Driven Executive Application User's Guide," PASC Tech. Report ZZ20-6438 (Jul 1978).

27. C. R. Hollander and H. F. Smith, "5100 Basic Compiler Prototype," PASC Tech. Report ZZ20-6444, (Oct 1978).

28. A. Hassitt, L. E. Lyon, "Array Theory in an APL Environment," APL79 Conference Proceedings, Rochester, (May 1979).

29. R. G. Scarborough and H. G. Kolsky, "Improved Optimization of FORTRAN Object Programs," IBM J. of Res. and Develop., vol. 24, No. 6, pp. 660-676 (Nov 1980).

30. P. A. Smith, A. Hassitt, N. R. Mendelsohn, P. Sih, "Experimental Pascal Interpreter for the IBM Series/1," PASC Tech. Report ZZ20-6451, (Feb 1981).

31. IBM Installed User Program 5798-DFH, "FORTRAN Utilities for VM/370," (1981).

32. L. M. Breed and H. F. Smith, "A Pascal Compiler for Series/1," PASC Tech. Report ZZ20-6453, (Sep 1981).

33. A. Hassitt and L. E. Lyon, "Pascal/DW: Pascal for the IBM Displaywriter," PASC Tech. Report ZZ20-6457, (May 1982).

34. R. A. Blaine, "IBM FORTRAN Overview, Portability and Conversion," PASC Tech. Report G320-3440 (Nov 1982).

35. A. Hassitt and L. E. Lyon, "Pascal/PC: Pascal for the IBM Personal Computer," PASC Tech. Report ZZ20-6461, (Nov 1982).

36. J. Littman, "The First Portable Computer," PC World, Vol. 1, No. 7, pp. 294-300 (Oct 1983).

37. S. M. Roberts, H. J. Myers and H. F. Smith, "Compound AND and OR Conditions in Decision Tables," submitted for publication (1984).

38. A. A. Dubrulle, "The Design of Matrix Algorithms for FORTRAN and Virtual Storage," PASC Tech. Report G320-3396 (Nov 1979).

Name Internal Zip City & State

HISTORIC

To: Mr. J.A. Haddad CHQ Valhalla

IBM

From Date: November 3, 1977

Name & Tie/Ext.: H.G. Kolsky 8/624/3136

Title/Dept. Name: IBM Fellow 60A

Internal Zip/City, State: P.O. Box 10500

or U.S. mail address: Palo Alto, CA 94304

HISTORY

~~IBM CONFIDENTIAL~~

Subject: Origin of the domestic Scientific Centers

Reference: Your request for information

The idea of establishing the domestic Scientific Centers came from a series of discussions and informal planning meetings in late 1963 and early 1964. This was also the time when Project X and Project Y (which later became the Mod 91 and ACS) were being started. The first reference I could find in my notes to the idea of setting up a separate group was in a meeting with a member of Charley DeCarlo's staff on Nov. 27, 1963. The other topics discussed at that meeting could well serve as an agenda for a Nov. 1977 meeting, if the names were changed: new product plans, competition, the need to study major customer problems, the importance of continuing professional contacts with scientific customers, and the need to provide input to the hardware developers of the needs of such customers.

Needs

February 10-14, 1964,a Planning Conference for Systems Research and Development (the original name of the Scientific Centers) was held in Los Angeles. A copy of this report is enclosed. Note that the flavor of the document is not one of pleading to establish the centers, but rather one of taking the lead in solving a need which was widely recognized by Company management. Also note that the same problems exist today.

The Centers were transferred from Corporate to the DP division on March 25, 1965. A copy of the bulletin board announcement is enclosed. Before this decision was made there had been a number of discussions as to where the Centers should report. Charley DeCarlo always said that they should be in a line organization in order to take the leadership role he envisioned for them. The alternatives seriously considered at the time were DPD and Research. Mr. Frank Cary (then president of DP) wanted the centers badly and "out bid" the Research division. He also offered to take Herman Goldstine from Research to lead the new organization. At the pre-announcement meeting held at Yorktown on March 23, 1965, a whole group of executives spoke of their hopes and best wishes for the centers:

E.R. Piore said the centers were to be congratulated for their success and stated that their work would continue and prosper.

IBM CONFIDENTIAL

C. DeCarlo reviewed the mission of the centers and stated that the time had come for them to move from staff to line. He viewed them as a new way of doing business.

F.T. Cary spoke of the Centers' long term technical sense of direction and their importance to DP.

G.B. Beitzel, then VP of DP market operations, described how the centers would fit in his organization. He then spoke of the importance of the maintaining confidence and respect of the Universities. We should do new applications, make changes to products, etc. not just serve as a means of information flow. We should bring in good people into DP. Joint studies are a "fine thing". Market requirements, case studies needs were expanding. The centers should do long term work and provide direction to DP. Finally he stated, "The Centers' mission will be protected".

A.G. Anderson, then head of Research, spoke of joint studies between the Centers and Research. He visualized a continuing interchange of individuals having unique skills on projects of mutual interest. Jack Bertram, Yorktown, was to be the contact point.

The Centers continued on their original path with ever-increasing resources for about 3 years. In 1967-68, however, university contacts and science in general started to fall from favor. (This trend has only recently reversed itself.) In the mean time the Scientific Centers have eroded to about half of their peak strength. In addition about 10 percent of the remaining man-power counted is in scientific marketing not in the Centers.

If you want more information please give me a call.

H.G. Kolsky

cc: Dr. L.M. Branscomb
Dr. L. Robinson

HGK:gm

Enclosures

HISTORICAL DOCUMENT H G Kolsky

Save, this the original planning conference for SR&D

Dp PAC – 3

REPORT ON PLANNING CONFERENCE

SYSTEMS RESEARCH AND DEVELOPMENT CENTER

AT LOS ANGELES

FEBRUARY 10-14, 1964

PLANNING CONFERENCE

Mr. I. Boldt — prog. devel. LA
Mr. B. M. Brown – Wash. DC
Dr. B. Dimsdale – LA
Dr. H. H. Givin, Jr. – LA
Dr. D. Heggie – LA
Dr. S. Jamison – LA
Dr. J. Mount – LA
Dr. D. V. Newton – WP

SYSTEMS RESEARCH AND DEVELOPMENT DEPARTMENT

IBM CONFIDENTIAL

IBM CONFIDENTIAL

INTRODUCTION

The purpose of this report is to summarize the major points of discussion and conclusions derived from the planning conference held at the Systems Research and Development Center at Los Angeles, February 10-14, 1964. Organized in three sections, it first discusses the current state of the scientific market; second, it proposes a general SR&D plan for supporting the scientific market through the SR&D centers, and finally, it outlines suggested operating policies for each of the centers.

SECTION I

CURRENT STATE OF THE IBM SCIENTIFIC MARKET

A characteristic of IBM customers who apply data processing equipment to scientific problems is that of rapid change. The very nature of the applications demands change, advancement, as a measure of progress. Technological and theoretical advances frequently result in an alteration in the use of and demands placed upon general purpose computing systems. Equipment which may be well suited to the customer's environment when it is installed, can suddenly be regarded as a limitation when viewed in the light of a technological or theoretical advancement in the application area concerned. Response to rapidly changing requirements thus becomes the keynote for SR&D.

1.1 Customer Environment

1.1.1 Professional relationships within the customer's business

Computing has become a major economic item for IBM's scientific customers. With this increase in significance has come strained professional relationships within many customer organizations which work to the disadvantage of the IBM marketing effort.

Specifically, competition between our scientific customers has led them to seek methods of minimizing the unit cost of computing many times at the expense of not providing adequate problem "turnaround"or"job" time service to the professional problem solver. Rather than installing multiple small computers, most customers have moved to larger systems because of the improved price/performance ratio available, as their computing volume increases. Even when IBM's largest processor alone could not satisfy the computing requirement of a customer, the centralized approach was pursued on the basis of economy in operations and physical plant even though the unit cost of computer performance is not lowered when additional like processors are installed centrally. Again prompted by the desire to lower the unit cost of computing, centralized computing facilities tend to implement loading rules that favor efficient use of the large computing system somewhat at the expense of user satisfaction. (Throughput is obtained by trade off against job time.) This situation gives rise to the polarization of the customer organization into two groups. The central computing authority is provided sole jurisdiction over computer acquisition and operation by virtue of economics; and the application oriented user group who is responsible for the productive work and who feels his needs are not being met in terms of his required job time.

The disadvantage afforded to IBM marketing teams by the polarization described is two-fold. In the order of increasing significance, the first aspect is that some applications simply are not getting programmed because of processing delay. Programmer productivity is more a function of the number of man-machine interactions provided that it is of total machine time made available. Excessive job time means less frequent interations between programmer and computer and detracts materially from realizing full application potential. The second disadvantage is over restriction of IBM sales activity in application areas which lie outside the domain of the central computing authority. IBM representatives can be restrained from assisting a remote user who wishes to acquire a free standing computer. Open threats to the centrally installed systems suffice to curtail IBM acitivty in outlying areas. On the other hand, competitive representatives are not so constrained. Further, to the application oriented user who is contemplating the acquisition of a computer "IBM" is synonymous with the central authority. A non-IBM system then would have the added feature of implied immunity from centralized control.

1.1.2 IBM Image

In addition to the above considerations which arise from IBM being the prime supplier of equipment installed at a central computing facility, the customers environment also reflects a general regard for IBM that stems from the company's size and rapid growth. In particular, the technical competence within IBM is not entirely evident in the IBM field operation. The pattern seen by the customers is for capable but inexperienced technical people to be assigned to accounts, quickly gain experience and soon move to more responsible positions. The procedure is necessary to satisfy growth requirements but it does little to enhance the customer's regard for IBM's technical competence.

An additional point relative to the "size" of IBM as seen by the customer relates to responsiveness. IBM presently is and will continue to be basically oriented to the broad requirements of the market. This requires that compromises be struck in hardware design, programming support, etc., so that IBM can satisfy the specific needs of most of its customers. On the other hand, the scientific customer must be imaginative, resourceful, and technically competent by the nature of his endeavor. As a result, IBM's effort to satisfy the needs of the "average market" is viewed by the scientific customer as a lack of response to his requests and suggestions.

1.2 Present IBM Approach

1.2.1 Marketing

For the most part scientific customers have moved rapidly from an era of dependence upon IBM for guidance in computer application to one in which a significant technical competence in various disiplines is required to keep abreast of the advances being made in even a single geographic locality. Normal expense/revenue objectives within the IBM branches do not encourage either the professional specialization nor long range pursuit of the activities in which prestige scientific customers are engaged. As a result, a significant loss in account control has occurred. Whereas the opinions of IBM representatives were previously regarded as coming from a source thoroughly qualified, the feeling has given way to apprehension. Since professional acceptance has diminished, any firm recommendation on the part of IBM can easily be regarded as biased self interest. The undertone of technical competence is lacking.

1.2.2 Product Development

A further observation can be made regarding IBM's present approach to the scientific market and this lies in the field of product development. Consider the unit record approach to scientific applications. Involved in the operation are multiple units (sorter, tab, calculator, etc.), each is separately programmed in that operators are instructed how each machine is used for a particular application. Also there is significant parallel data flow, e.g., data are being sorted concurrently with the calculation on other data.

Product development, specifically computing system development, is leading to analogous operation. Separately programmed units, e.g., data channels of the 709X, and the IBM 704X of the directly coupled 704X/709X provides asychronous data manipulation within IBM's current processor line. Particularly in the more advanced installations where the number of programmable units increases as in the case of the IBM 7094/7040/7740/1301/1311, the question of system control arises. In the unit record approach central control is manifest in the supervisor who implements loading rules which strive toward maximum performance within the operation. Unlike the role of the unit record operator which has been replaced by hardware counterparts (data buses for information transfer, sequentially programmed operating sequences, etc.), the role of the supervisor is relatively undeveloped. It is a theoretical consideration, being a program, rather than a hardware device. It must audit several units concurrently (multiprocessing) and it must employe explicit algorithms which balance efficiency against response time for the system. In brief, decentralized systems control, by default is leading to unregulated job times in the operation of large scale scientific computing facilities:

opcu syst

1.2.3 Research

Acknowledging that the central control function of a scientific computing system is a programmed concept rather than a hardware device leads one to consider system control techniques as a topic of basic research. Presently, we have technology leading research. Solid state advances serve as an example of this situation. On the other hand, research in programming the control of highly parallel systems is extremely limited.

1.2.4 Application Development

In the existing approach used by IBM in the area of scientific application development there are a number of examples of work which in themselves are accomplishments. The necessary, but lacking, factor which would bring the application development to fruition is a combination of systems support and professional contacts through which the work could be introduced in the community.

More specifically, SR&D's contribution in the area of application development will be the development of prototype systems as opposed to officially supported application programs. Once the level of refinement is determined by SR&D through mutual evaluation with the scientific customer, such documentation and systems refinements as necessary will be provided through normal channels.

Where the importance of the customer and the nature of the application warrants SR&D may look to the development of complete integrated application systems such as user oriented language systems or any other programming support which makes the computer more accessible to the scientific user.

1.3 Systems Technology

There are two important indicators in systems technology that suggest the direction in which scientific computing may move.

1.3.1 Fundamental Constraint

One indication arises from work being done toward the design of the next generation large scale systems. Clearly and simply, pulse propagation speed is a limitation. Distances deemed short by electronic packaging standards are long when converted to the propagation delays they introduce to confound the digital engineer. As a result one can conclude that future systems will contain more parallel organization wherein processes will be performed on the basis of the parallelism inherent in their structure. Little is known about how to exploit parallel system capability effectively at the application level. The previous challenge of automatic memory allocation, for example, would be superceded by questions of multiprocessing such as automatic unit allocation.

1.3.2 Communications

The second indicator to be considered lies in data communications. Progress in instrumentation, telemetry, and data transmission shows that technology exists whereby remote scientific computing can be implemented when the concept is proven. The possibility is strong that within remote scientific computing lies the principles governing the trade off between throughput and job time in the scientific computing facility. Irrespective of the development of remote scientific computing the trend toward remote data acquisition is clear. The future holds the answer as to what interdependence exists between remote scientific computing and real time or on line data acquisition. Whatever the relation the evolution will be rapid.

SECTION 2

GENERAL SR&D STRATEGY

2.1 Evolution in the Scientific Market

The point was made in Section 1 that rapid change is a characteristic of the scientific market. Generally stated, then the function of SR&D is to provide within IBM the ability to match rapid change with correspondingly rapid response. The primary mode in which this function will be performed is by establishing long range contacts with leading scientific customers, understanding their problems, defining solutions, and insuring that capable IBM agencies respond to the need.

Within SR&D the allocation of resources will be performed on the basis of rate of change. Those scientific areas in which maximum rate of change is either observed or anticipated will receive the greatest emphasis in SR&D activity.

2.2 Marketing Relationships

SR&D operations will be consistent with IBM policy in that activity within accounts will be conducted with the cognizance of DP representatives concerned. In particular, a close operating relationship is anticipated with Scientific Marketing - DP. For the most part, SR&D will look to the long range needs of the market which lie beyond the normal objectives of DP Sales and Scientific Marketing. Through long standing relations with key customers, it is anticipated that guidance can be rendered to IBM efforts which range from a local sales strategy to special installation assistance. By virtue of the insight gained by SR&D, IBM should be in a better position to introduce timely products in an effective manner.

Within the framework set forth in Section 1.2.1, SR&D activity will investigate remote scientific computing as a means of providing a satisfactory trade off between throughput and job time in scientific computing facilities. It is accepted that replication of installed systems is an immediate answer to customer growth because it is a fairly acceptable solution relative to the internal customer political situation, it is easily implemented by the customer, and it serves as an immediate holding action on the part of IBM. Looking ahead, however, the concept of compatible satellite computers will be explored concurrently with investigation of the remote scientific computing. Several factors foster concurrent satellite/remote scientific investigation. The first is that the satellite approach is available to the field through FORTRAN compatibility now and will be even more attractive when future products are announced. On the other hand, remote scientific requires some system development lead time but appears, at this point, to be the most promising approach because a) it is compatible with the centralized computing authority objectives, b) it reduces turnaround time to the order of minutes if not seconds as in the case of the conversational mode, c) printed output volumes would be reduced and, finally, d) remote scientific may prove to have economic advantages over the satellite computer installation.

Developing the remote scientific concept will require advances systems engineering technique which can serve as resources in themselves. Not only will SR&D system engineers have the knowledge and experience to prepare proposals for and assist in the installation of remote scientific systems, but through publications such as the IBM Systems Journal, can communicate their understanding throughout IBM.

2.3 Application Areas

Application development holds different benefits depending upon one's point of view. To the Data Processing Division application development provides a marketing tool i.e., current work by IBM in areas of business which are pertinent to particular sales situations. Also through application development effort DP personnel are trained in specialized fields.

To the Data Systems and General Products Divisions application development is a product development activity. In addition to supporting present products, application development provides the opportunity to determine future hardware design requirements.

In SR&D application development will mean being at the right place at the right time so that IBM can take advantage of those areas of science which are likely to evolve rapidly. The degree of participation will vary from a minor investment in some areas where it is desired to only have sensors established which note the direction of rather slow changes, to an active participation, on the other hand, in areas where the field is developing rapidly and it is important that IBM be prepared to respond equally quickly.

A number of specific scientific application areas of interest to SR&D are the following:

Numerical Analysis - While the fundamental principles of numerical analysis remain somewhat invariant, the algorithms by which the fundamentals are implemented reflect a specific machine organization. In this light, then, SR&D will be interested in the influence of multiprocessing on numerical analysis algorithms.

High Energy Physics - This field is an example of one whose computational needs are changing rapidly and it behooves IBM to participate actively in the area so that it can be prepared to answer the needs of these prestige scientific customers. Bubble chamber data reduction, for example, has moved from an off-line manual operation to its present status, an on-line semi-automatic data acquisition operation, in something less than three years.

There are several other aspects of high energy physics that foster SR&D interest. Both bubble chamber and spark chamber experiments result in large quantities of photographically recorded data. The physical meaning of this information is available only through significant computation on digital data obtained from the film by means which potentially can be further automated through curve following and pattern recognition techniques. Further, the data, in digital form, is cumulative and will draw heavily upon future computing capacity for subsequent re-analysis.

Nuclear Design - SR&D's interest in maintaining contacts in the nuclear design field is two-fold. Since nuclear design problems continue to press the largest and fastest computers available, the first, and obvious function is to provide IBM with an objective evaluation of the large machine market. Second, questions regarding the degree to which the large machine can satisfy general purpose requirements through software support can often be answered only through experience on installed equipment. SR&D professional contacts at the nuclear installation will provide for this type of feed back.

Space Technology - Little need be said to describe the rate of change in the area of space technology. The field is moving extremely rapidly and warrants heavy SR&D participation in both scientific theory, e.g., orbit calculations, and in the scientific systems considerations which are incorporated in space flight observation systems.

Medical - The rapid rate at which developments are being made in the area of medical computing as well as the financial support of these endeavors prompts SR&D interest in the field. While some market resistance exists at the clinical level, medical researchers are rapidly combining the talents of the life scientist, the instrumentation engineer and the mathematician into teams that place heavy demands upon computers.

Statistics - Rate of change in the field of statistics, particularly as applied to the behavioral sciences, suggests light SR&D attention while maintaining an awareness of total financial support being made available to the behavioral sciences.

Education - While public education is comparatively slow moving, cost motivation is stimulating the area of industrial education into rapid progress. SR&D will look then to industrial education as an application area of interest.

Geophysics - Geophysics together with the associated fields of oceanography, seismology, and meteorology is an application area in which increasing interest in the scientific community is accelerating change in the computing requirements. Applications within this field have the additional features of large quantities of data, e.g., TIROS photographic data, and the numeric solution of three dimensional differential equations with the associated demands placed upon computing systems.

2.4 Systems Orientation

SR&D, as the name implies, will have a systems orientation. With this orientation SR&D will be in a position to complement DP activity aimed at NSRI objectives and DS/GP profit motivated efforts by providing specialized capability in organizing IBM products into systems which meet the needs of the scientific customer. For the most part, SR&D support of DP activity will consist of showing the scientific community outstanding systems in operation on tasks set forth by competent professionals to demonstrate the system in a meaningful manner and consult regarding special needs. In ~~a number of~~ situations where the system is unusually complex, the SR&D Center may accept the responsibilities of providing ~~the programming and~~ systems ~~engineering~~ support.

SECTION 3

OPERATIONS

3.1 Headquarters Operations

Headquarters operations will be limited to overall technical planning and administrative services. In particular, SR&D headquarters operations will coordinate formal communications and distributions between centers and other parts of IBM. One such coordinating function will be liaison with the IBM Systems Journal: a second will be the maintenance of a central listing of problem areas within the scientific market that must be solved by assignment within SR&D or elsewhere within IBM.

3.2 SR&D Field Operations

3.2.1 Organization

The major portion of SR&D activity will consist of operations in and about approximately seven SR&D Centers located in areas of concentrated scientific endeavor. Each center will be staffed with a number of professionals in disciplines commensurate with the objectives of the center.

Since the SR&D Centers may be oriented to the laboratory environment, such as Los Gatos, or about a major system, such as the 7040/7094 WDPC system, the relative distribution of staff by discipline will vary between centers. In either event two types of professionals will be assigned to the centers - systems specialists and application specialists.

3.2.2 Field Center Operating Policy

Several policies will govern operations of all SR&D Centers.

Response

Since the purpose of SR&D is to establish a means by which IBM can match its actions to rapid change within the scientific community, responsiveness must be the basis for SR&D operation. In providing timelv response to the scientific customer's needs, SR&D will first look to current projects within such IBM agencies as are capable of providing the service. If, however, it is found that IBM has not anticipated the need and the requirement significantly reflects upon IBM's image in the scientific market, SR&D will either recruit support of the IBM divisions, if possible, or fill the requirement with its own resources.

Professionalism

The aim of SR&D is to establish contact with scientific customers on a professional basis. In keeping with this objective SR&D Center operation should provide technical persons with an opportunity to further their professional stature through publication, participation in professional societies, and through work assignment itself.

Esprit de Corps

Rapid response and maintenance of a professional approach tend to be mutually exclusive. Through a keen esprit de corps, however, SR&D. professionals can be provided with an atmosphere conducive to the extra effort required to cover the conflict between response and professionalism.

Recruiting and Training

A function of the SR&D Centers is to recruit professionals, train them in computer technology, if necessary, and make them available through transfer to the line divisions. In the operation of SR&D centers the need to inject a degree of professionalism into other parts of IBM should be borne in mind.

Range of Operation

Operations of SR&D centers should be geared to the long range contact with scientific customers. DP through local sales coverage supplemented by Scientific Marketing is equipped and motivated to handle the immediate short range association with customers. SR&D, however, should work toward the long standing professional relations which can continue through fluctuations in the local sales atmosphers. It is clear, then, that established professional standing should not be compromised for short term reasons.

3.3 Functional Assignments

3.3.1 LOS ANGELES CENTER

Director: Dr. H. H. Givin, Jr.

Location: WDPC at UCLA and Kirkeby Center
Los Angeles, California

Application Specialties: Space Technology

Medical

Industrial Education

Geophysics

Systems Specialty: 7740/7040/7094 Direct Couple Remote Scientif

Account Affiliation:

CIT/JPL	Space
UCLA	Medical/Geophysics
SDC	Industrial Education

3.3.2 SAN JOSE CENTER

Director: J. W. Luke

Location: Los Gatos, ASDD Laboratory

Application Specialty: High Energy Physics

Nuclear Design

Terminal oriented applications

Automated instruction

Systems Specialty: Terminals

Conversational mode remote scientific

Account Affiliation: LRL – Livermore Berkeley

Stanford Linear Accelerator (SLAC)

DP 9AC – 3

3.3.3 WASHINGTON, D.C. CENTER

Director:	Unannounced
Location:	Washington, D.C.
Account Assignments:	Federal

3.3.4 HOUSTON, TEXAS

3.3.5 CAMBRIDGE, MASSACHUSETTS

Director:	Norman Rasmussen
Location:	Cambridge, Mass.
Application Specialty:	Engineering Design
	Automated Instruction
	Behavioral Sciences
Systems Specialty:	Time sharing
Account Affiliation:	MIT

3.3.6 CHICAGO CENTER

1968 Statement of Scientific Center Mission

MISSION

The Scientific Center mission is fourfold:

<u>Communication</u>: A vital part of Scientific Center mission is to maintain an effective link between various technological entities both within and outside IBM. Such entities would include Industry Development and Marketing, Universities, Advanced Technological Customers, Product Marketing, IBM software and hardware design group, and other IBM Research and Advanced Technology activities.

Interact with Sci. Community

<u>Initiation</u>: Scientific Centers are also directed towards identifying advanced technological application and system areas whose development would broaden the computer marketplace. Scientific Centers would initiate projects and execute feasibility studies to give initial impetus to these new areas of endeavor.

do Ad Tech

<u>Propagation</u>: Scientific Centers will work, primarily with Industry Development groups, to provide technical consultation and transition of advanced development activities. In order to complement this aspect of the mission, Scientific Centers will also work, primarily with SDD, to assist in initiating appropriate hardware and software development activities necessary to maintain and increase IBM's position in the marketplace of the future. They will also propagate their activities through education, demonstration and other modes of dissemination.

Interact with Development

<u>Image</u>: Scientific Centers will maintain high image in the outside technological community in order to be focal points for attracting high-caliber technologists to DP Division and the IBM Corporation as a whole.

Extern. Imag

IBM Corporate News Bulletin

SAVE: HISTORICAL DOCUMENT SR&D goes from Corporate Hq to D

I. ?. Kolsky
~~San Jose~~

March 25, 1965

Db PAC - 3

SR & D ACTIVITIES REALIGNED

The functions of the Systems Research and Development staff have been realigned to strengthen the link between IBM's industry-oriented market planning and the leadership institutions in the scientific and academic community.

Responsibility for the SR & D centers and the Mathematics and Applications Department have been assigned to the Data Processing Division. The centers are located in Cambridge, Mass., Houston, Texas, Los Angeles, Calif., Palo Alto, Calif. The mathematics and Applications Department is located in New York City. Dr. Herman Goldstine, formerly director of mathematical sciences for the Research Division, has been promoted to director of scientific development and will head these operations for DPD.

The Systems Journal becomes a responsibility of the office of Jerrier A. Haddad, IBM director of technology and engineering.

Dr. Charles DeCarlo will continue to report to Arthur K. Watson, IBM senior vice president. In addition to supervision of the Systems Research Institute, Dr. DeCarlo will take on a special assignment in the area of large scale scientific systems.

HISTORIC   E/ME/A Scientific Center FEB 18 1975
CHARTER

PROF. DR. KARL GANZHORN
DIRECTOR OF SCIENCE & TECHNOLOGY IN EUROPE

IBM
7000 STUTTGART 80
POSTFACH 800 880
PASCALSTRASSE 100
TELEFON (0711) *785-4001

February 14, 1975

Dr.
H.G. Kolsky
IBM Fellow
IBM Corporation
Armonk, N.Y.
USA

Subject : Scientific Review Board

Dear Mr. Kolsky,

During the past two years our IBM Scientific Centers have reassessed their role, functions and projects in view of the changing scientific and business environment.

As a result a

" SCIENTIFIC CENTER CHARTER "

has been established and agreed upon between IBM Europe and the Country General Managers (attachment 1).

In order to support the Scientific Centers further in executing their missions and tasks a

" SCIENTIFIC REVIEW BOARD "

of IBM Europe will be formed (attachment 2). It is a pleasure for me to invite you to become a member of this Board and I would be pleased if you would accept this senior function in order to make your professional insight and experience available for the benefit of the Scientific Centers.

The objectives of the Board are to review and to assess the programs undertaken in Scientific Centers as outlined in the Board Objectives (attachment 3). We want to keep an option of extending the Board's advisory function to other science-oriented issues within IBM Europe if appropriate.

The Board will meet two times two days per year, preferrably at one of the Scientific Center locations.

The first Board Meeting will take place on April 22, and 23, 1975 in Madrid. A separate invitation will follow.

Sincerely yours,

K. Ganzhorn

Attachments :
1. Scientific Center Charter
2. Scientific Review Board Member List
3. Board Objectives

cc.:
see distribution list

IBM EUROPE

SCIENTIFIC REVIEW BOARD

MEMBERS

Chairman :	Prof. Dr. K. GANZHORN Director Science and Technology in Europe, Stuttgart
Secretary :	Dr. G.S. KOZAK Science and Contribution Programs Director, Paris
Members :	J.W. FAIRCLOUGH Managing Director Development Laboratory, Hursley
	F. GENUYS Product Line Management-Programming, Paris
	Dr. H.G. KOLSKY IBM Fellow, Armonk
	C. MAZZESCHI Director of International Programs, Milan
	P.W.L. MORGAN Group Director DP Marketing, Paris
	Prof. Dr. W.E. PROEBSTER Director of Research and Development Coordination, Stuttgart
	Prof. Dr. H. ZEMANEK Director of SDD Laboratory, Vienna

IBM EUROPE

SCIENTIFIC REVIEW BOARD

OBJECTIVES

The Scientific Review Board of IBM Europe will be established in order to review, assess and support the programs and operations of Scientific Centers .

The objectives of the Board are :

1. To undertake periodic reviews of the Scientific Center activities, studies and plans, approximately once per year and center with particular emphasis on

   a) meaningful orientation of missions, projects and future plans according to the Scientific Center Charter,

   b) scientific and technical excellence,

   c) proper coordination and linkage of the programs with other IBM functions and with potential or identified interested parties.

2. To assess the major objectives and management aspects of Scientific Centers together with local line management.

3. To advise local and IBM Europe management on issues with powers reserved to IBM Europe S.A. relating to the establishment of Scientific Centers and joint projects with outside partners.

Distribution list of carbon copies :
(without attachment 1)

F. de Asua, Madrid
W.A. Boesenberg, Stuttgart
L.M. Branscomb, Armonk
A. Cacciavillani, Milan
K. Cassani, Paris
B.C. Christensen, White Plains
P.C. Clarke, Paris
D.Z. Cohen, Tel Aviv
F.J. Cummiskey, Paris
J.A. Haddad, Armonk
L.F. Hanau, Vienna
H. Haussmann, Paris
J. Herbart, Paris
J.G. Maisonrouge, Paris
E.R. Nixon, London
G. Sandorfi, Paris
O. Weideli, Paris

H. G. Kolsky
FEB 18 1975

IBM EUROPE

SCIENTIFIC CENTER CHARTER

FEBRUARY 1975

CONTENTS

A) ENVIRONMENT

In the changing environment of the 1970ies, IBM's Scientific Centers can become a novel type of industrial Research & Development resource which (besides the classical industrial objectives) can specifically address our responsibilities extended towards society aspects.

The overall scientific and technological interface of our company towards the outside world has to take into account three areas (see chart 1) :

1. The industrial community, which is based on the principles of free trade, competition and profit.

2. The scientific and academic community with its two fundamental objectives, namely :

   a) to enhance knowledge of men through pioneering research and thinking

   b) to promote the scientific foundation of all facets of cultural, human, spiritual and economic life through oriented research and exploration.

3. The general society, which must address economic, cultural, educational and ecological evolutions.

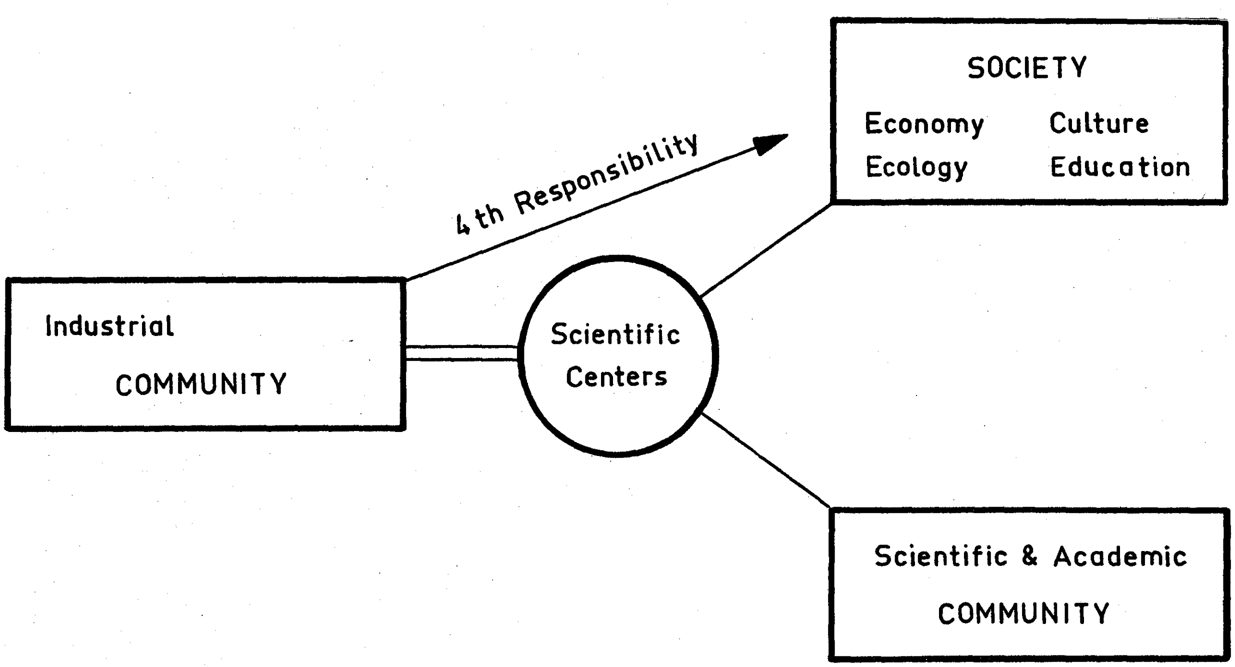

Feb 14, 1975

B) MISSION

The Scientific Centers will select and concentrate on one or more fields in which they will strive for continued first-class scientific competence. They will aim in three directions :

1. **Applications Research**

   Within their fields, the Scientific Centers are directed towards identification and exploration of advanced applications, software technology and system areas, whose development would have a scientific or social importance and, potentially, could broaden computer use. Scientific Centers will initiate and perform projects as well as feasibility studies to promote these new areas and to develop new opportunities for IBM's contribution to society.

2. **Scientific Information Flow**

   Scientific Centers will act as one of IBM's interfaces to the external scientific community (example : technical and scientific groups, professional societies and universities) by providing consultation and two-way communication in their fields of competence.

   Results of project work done in partnership with an outside institution will generally be available for public use or for use by the participating institution. If, in the course of a study, a contribution to product development becomes foreseeable, then the study must be dissociated from this part of the work, which will have to be appropriately protected.

3. **Image**

   Scientific Centers will strive to achieve the respect and confidence of the outside scientific and technological community. In executing their charter, they also assist Company External Relations function nationally and internationally.

C) STRUCTURE

In order to fulfil its mission, the structure of a Scientific Center must consist of two types of manpower resources complementing each other :

1. A competence nucleus for each mission orientation. This nucleus must go beyond monitoring the status of the art in its area of competence. It must become a recognized partner in science through qualified publications etc., while at the same time, it remains in close touch with our industrial aspects. Unless such a group can be hired with established competence, we have to allow a build-up phase of 5-7 years, which is the time constant of a good scientist to establish himself in the forefront of science. Depending on the specific field, this competence nucleus can be anywhere between 5 and 15 scientists. With present resources, we should aim at limiting the number of competence nuclei to 1 or 2 in each Center location.

   An adequate flow of personnel, plus stimulating fellowships, must be undertaken continuously in order to maintain continued scientific vitality.

2. Surrounding this competence nucleus must be project teams which implement the Scientific Center objectives with projects carried out with outstanding technical excellence. The selection of mission concept and projects should comply with the criteria in Section D, page 5.

   The project teams can be composed of members of the competence nucleus, of fellowship assignees from the academic community, of assignees from other IBM functions, e.g. systems engineering, programming development or industry centers, development or research laboratories. The lead-time to visible and applicable results of such a project must be in the timeframe of 2-3 years with objectives explicitly determined at the beginning. Taking the basic objectives into account, the projects should be selected from country or international origins in such a way that they benefit direct from the competence nucleus with which they are associated. Yet it may well be that they also, in a supplementary way, have to rely on nuclei in other Centers through exchange of information and project assignees.

It is important that the project teams remain flexible in their manpower composition. Assignment to such a team from other IBM departments must become a distinction and challenge for high-class work. It must also have a mutually seeding effect for the departments involved.

SCIENTIFIC CENTER STRUCTURE

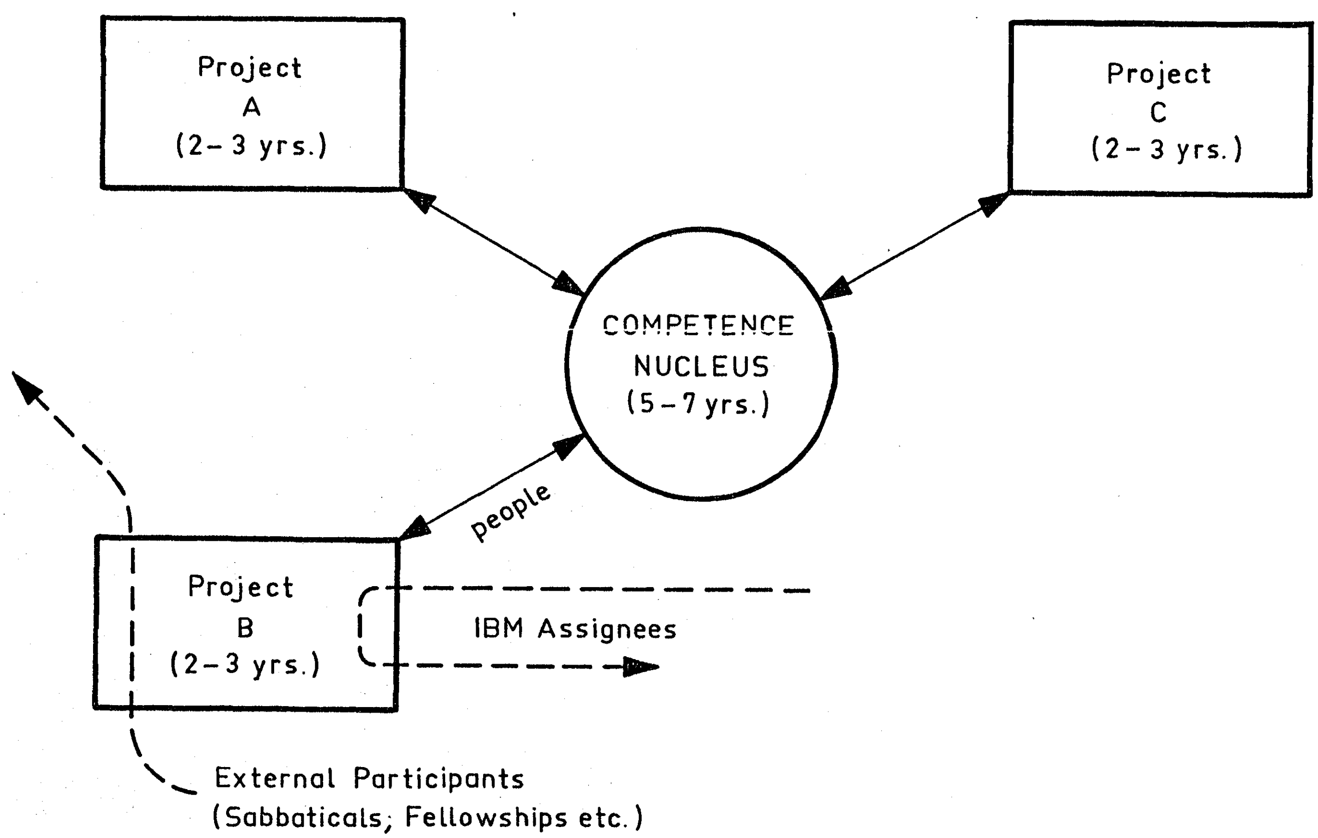

D) CRITERIA FOR MISSIONS AND PROJECTS

The establishment of long-range center missions and competence and short-range objectives should be guided by a majority of the following criteria in selecting fields and projects :

1. Are they at the leading edge of the art in science and industry ?
2. Do they explore new fields of computer usage offering short- or long-range benefits to IBM or to the society ?
3. Do they support IBM's role as a partner in society, nationally and/or internationally ?
4. Do they merge the technological and scientific progress of the industrial and scientific communities ?
5. Do they promote interaction between the scientific and industrial communities ?
6. Do they promote information flow between the scientific and industrial communities ?
7. Do they constitute a prime example of what scientific excellence, combined with industrial effectiveness can achieve, serving social and company interests ?

E) REPRESENTATION OF BROAD FIELDS

Society-oriented responsibilities may have to be oriented towards very broad fields, e.g.

" Environmental Problems "
" Ecology and Urban Problems "
" Computers in Medicine " etc.

With the present and foreseeable magnitude of our Scientific Centers, it is impossible to have the total necessary competence of such fields in any one center. Thus the competence nuclei plus individual projects of several centers together can only be representative of total Scientific Center competence and efforts in such broad fields.

In order to direct questions and channel requests from the outside to the various competence groups properly, the Scientific Center Managers in each country must serve as the prime contact and interface. Therefore, they have to maintain a current overall awareness on all Scientific Center competences and projects. In doing so, the company can be responsive towards external requests on a local and national basis.

F) PLANNING

In order to fulfil its mission, each Scientific Center should establish long-range missions and competence and short-range objectives, both to be included in a Scientific Center Strategic Plan. This plan should cover :

1. **Missions**

   Description of the fields of endeavor and their internal and external relevance.

2. **New Applications**

   Identification of the current fields of activity, the nature and scope of anticipated feasibility study areas and the nature and scope of major project areas.

3. **New Techniques - Software Systems Development**

   Identification of the current fields of activity, the nature and scope of anticipated feasibility study areas and the nature and scope of major project areas.

4. **Scientific Relations/Communications**

   Identification of the anticipated main activities and the key external contacts such as university education and research, Government entities, professional societies, important individuals etc.

5. **Manpower and Resources**

   Manpower and resources requirements and plans to achieve the objectives.

6. **DP Equipment**

   DP equipment requirements and plans to achieve the objectives.

G) MANAGEMENT

The WT Scientific Centers are hosted and funded by the respective national IBM companies. Therefore, the country general management ultimately approves all activities in the center, after all necessary functional concurrences are obtained. It then also ensures the necessary performance evaluation.

IBM Europe Management, advised by the Director of Science and Technology for Europe, will establish appropriate means of advising and assisting the Scientific Centers' country management in establishing, reviewing and assessing the Centers' missions, objectives and performance.

Establishment of new Scientific Centers requires IBM Europe's line management review and approval, as does the establishment of joint projects with outside partners and study/research contracts with educational institutions. All formal contacts with the outside community are to be structured in accordance with current WT instructions. (See WTC Instruction Manual, Section 0-20 General Instructions - Pages A1, B1, C1 and M1).

Reprinted from

IBM

Systems Journal

Vol 28, No 4, 1989

History and contributions of the IBM Scientific Centers

by H. G. Kolsky
R. A. MacKinnon

History and contributions of the IBM Scientific Centers

by H. G. Kolsky
R. A. MacKinnon

The IBM Scientific Centers are celebrating their twenty-fifth anniversary. These worldwide Centers are autonomous organizations that provide IBM with the ability to respond rapidly to the evolution of computer technology for IBM and for its scientific customers. During the past quarter century these Centers have provided technical leadership in almost every branch of computer science. Today, the 17 individual Centers continue to explore new technical areas and provide significant contributions. This paper has three parts: an introduction to the mission, scope, and history of the Centers; a description of each Center's charter, history, and accomplishments; and an extended list of selected publications for each Center.

The IBM Scientific Centers were chartered, and the first four established, in 1964 to provide IBM with the ability to respond rapidly to the scientific marketplace and to changes in technology. From the beginning the primary mode of operation has been establishing long-range contacts with leading scientific customers, understanding their problems, defining solutions, developing prototypes, and ensuring IBM's responsiveness to their needs.

The idea of having a small, entrepreneurial organization within a much larger company has appeared again and again in American industry. "Skunk works," "ad tech groups," and "back room projects" are a part of American corporate folklore. The IBM Scientific Centers were created to fill just such a role for IBM. The emphasis was and is on having an autonomous organization of highly-skilled professionals who can develop scientific solutions and applications at the leading edge of technological change. During the past quarter century there have been many individual accomplishments within the Centers that have had a large impact on IBM, not only in terms of the product line, but also in terms of upgraded scientific and technical quality for other projects within IBM, basic research, market support, and a generally enhanced relationship with scientific customers. Thus, technology transfer remains the Centers' key contribution for IBM.

The approach and initial Centers were so successful that the concept of the centers was adopted on a worldwide basis. There are now 17 operational IBM Scientific Centers. The Centers have contributed in many branches of computer science and especially in computer architecture, personal computing, numerical analysis, algorithms, data structures, operating systems, relational databases, languages, compilers, microcode, networks, artificial intelligence, and knowledge-based systems.

Contributions in science and computer applications developed by the Centers include basic physics, chemistry, mathematics, agriculture, oceanography, surface hydrology, management science, econometrics, operations research, and commercial packages. The Centers have been heavily involved in human issues, such as the medical applications of computers, speech analysis and synthesis, computer vision, and natural-language projects in Spanish, Catalan, Arabic, Italian, French, German, Chinese, and Japanese. On a worldwide basis there has been recurring involvement in addressing problems of national importance to dozens of countries, such as regional planning, food and oil production, pollution studies, mapping the Nile, and protecting Venice from the sea. It is also common for the Centers to have joint studies in progress with many universities and on all aspects of their work. Figure 1 illustrates some of these activities.

The original objective of the Centers was to have each one involved in the professional and academic environment of a nearby university—working with professors, involving students in projects, and contributing to the scientific milieu of the university. In practice each Center has also operated on a national and international basis, participating in projects with governments and remote partners. It was also an objective that the Centers influence IBM's technology direction and product offerings. Both objectives have been consistently achieved over the history of the Centers.

It was also an objective that the Centers influence IBM's technology direction and product offerings.

United States Centers

When the first IBM Scientific Centers were founded 25 years ago in the United States, they were built on already firmly established traditions. The IBM Research Division was well-established, as were several large product development laboratories. The Applied Science Department in New York City under Cuthbert Hurd had been very successful in working with scientific customers during IBM's early expansion into electronic computing in the mid1950s. The key models were the joint IBM and University of California at Los Angeles (UCLA) Western Data Processing Center (WDPC), located on the UCLA campus, and a similar Center near the Massachusetts Institute of Technology (MIT), both of which had been established in 1957. Much of the philosophy and mode of operation for the new Centers came directly from these two prototypes.

The founder of both the earlier centers and the IBM Scientific Centers was C. R. "Charlie" DeCarlo, who was on the staff of IBM senior executive Arthur K. Watson. DeCarlo, with Watson's full support, foresaw a continuing need for an IBM commitment to applied science. The result is the present international network of Centers.

Four Centers were established in the United States in January 1964, with two more planned. The Los Angeles (California), New York City (New York), and Cambridge (Massachusetts) Centers, under the leadership of Homer H. Givin, Theodore I. Peterson, and Norman L. Rasmussen, respectively, were drawn from existing or associated groups. An entirely new Center, headed by John W. Luke, was established at the Los Gatos Laboratory in California and soon moved to Palo Alto, adjacent to Stanford University, in 1965.

After DeCarlo, Herman H. Goldstine was the first director of the new organization. The Centers have had many leaders, but it is Louis Robinson who must be given much of the credit for guiding them during the critical years from 1975 to 1982. Throughout this period he articulately and steadfastly promoted the importance of science and the Centers for IBM and its customers.

World Trade Centers

The IBM World Trade Scientific Centers had their roots in earlier organizations, as did the Centers in the United States. In 1960 a small group of scientists in Paris, headed by Rene Moreau, began projects in scientific development for IBM France. In 1965 a small group of scientists from IBM Italy started working at the University of Pisa. However, it was not until 1967, when the IBM Grenoble Scientific Center opened, that Centers based on the model used in the United States made their appearance.

From 1967 to 1972, twelve new Centers were opened around the world. Around 1980, six more were established, and several earlier ones merged with others and were closed. A number of organizations were also derived from the Centers, such as an IBM-sponsored center at the Asian Institute of Technology in Thailand; the European Center for Scientific and Engineering Computing at the IBM Rome Sci-

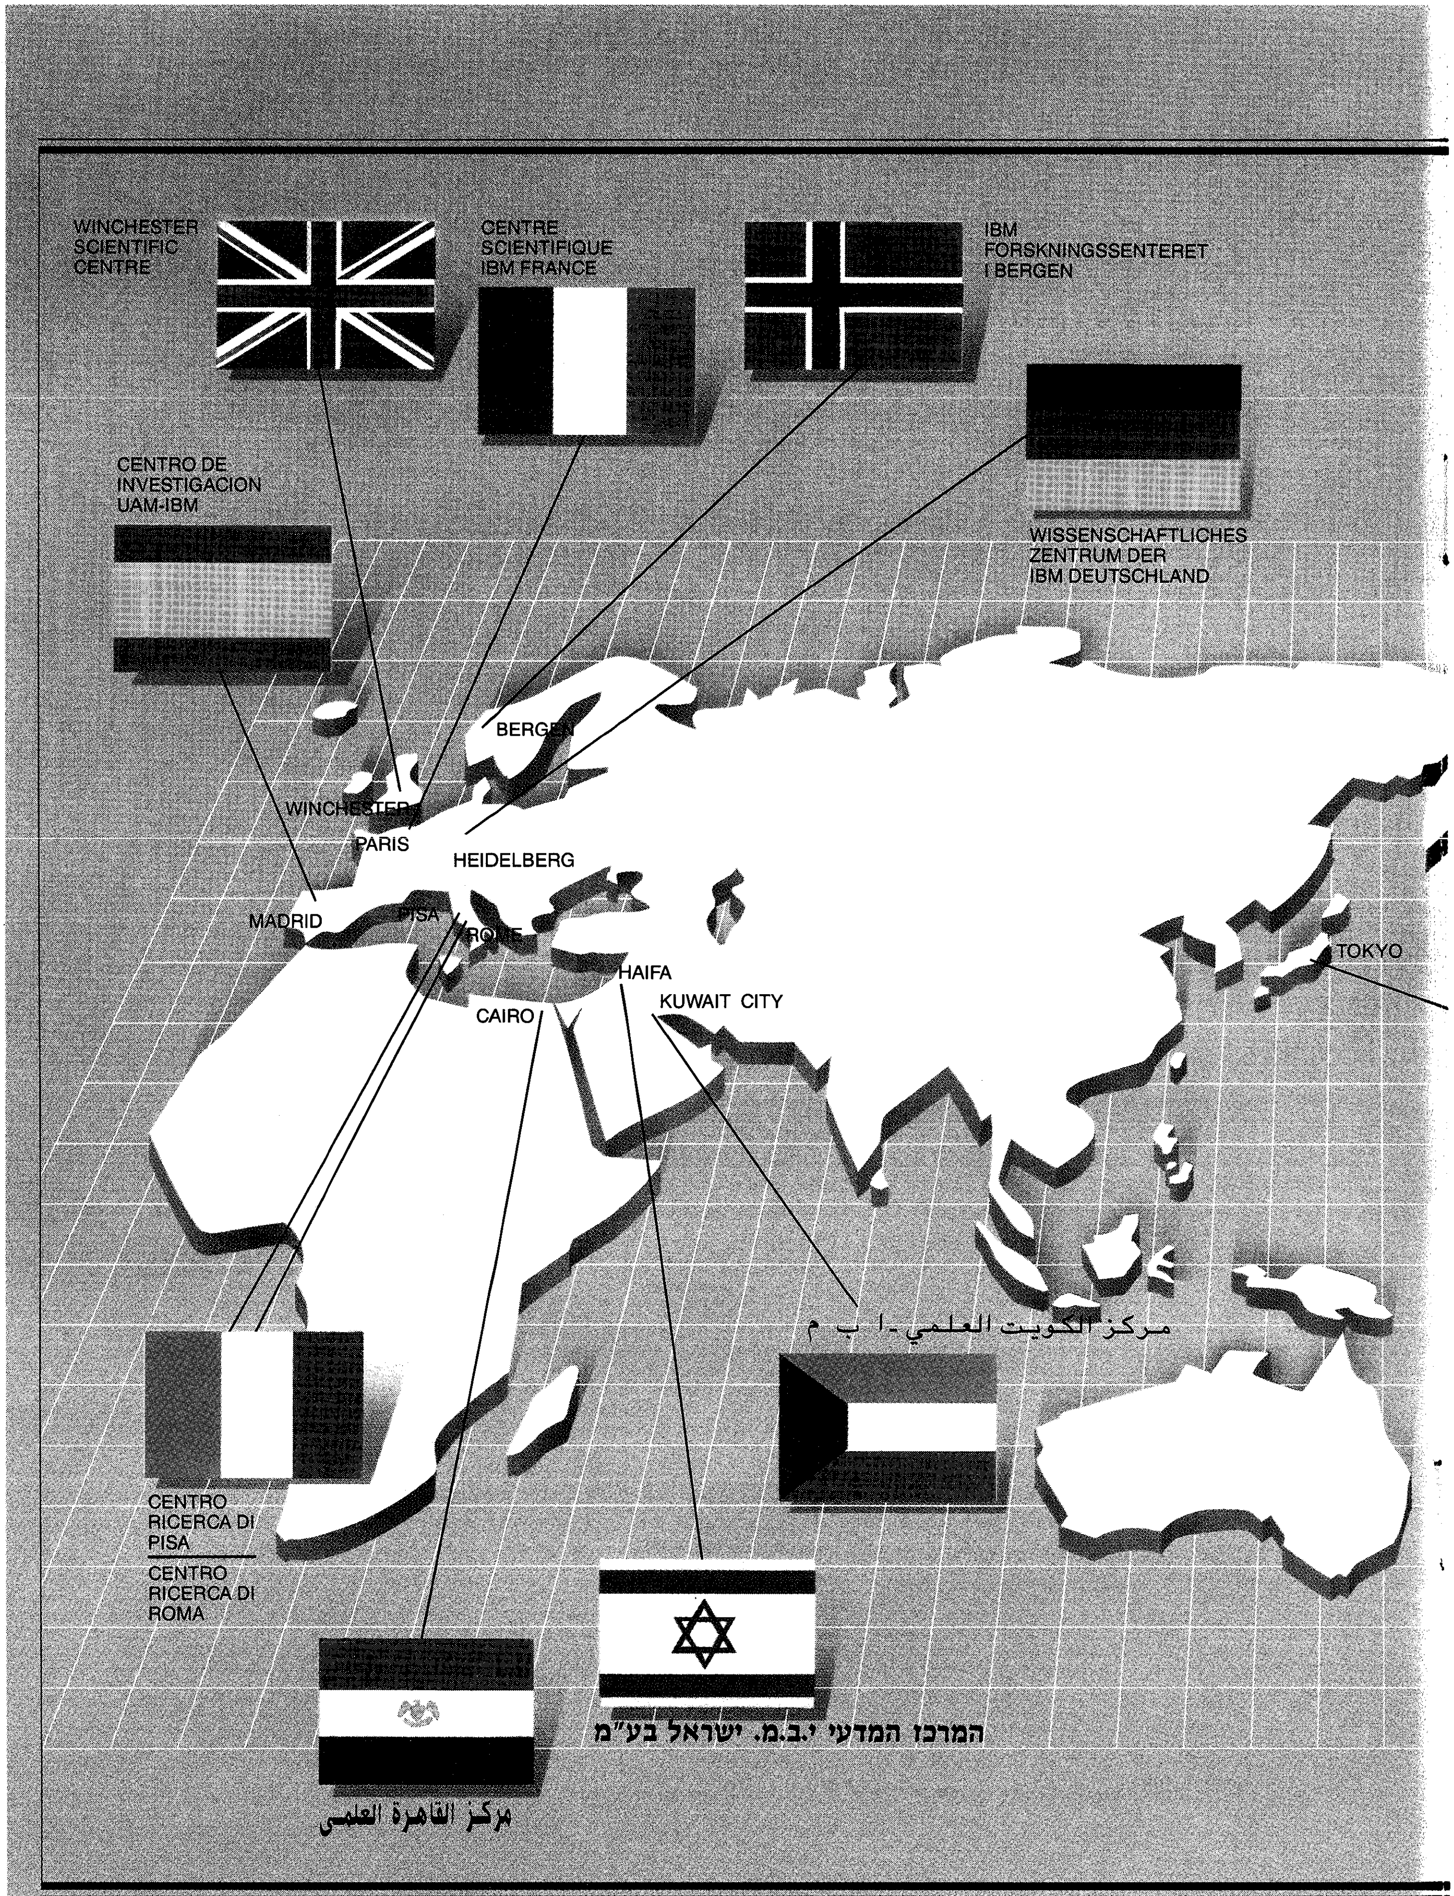

WINCHESTER SCIENTIFIC CENTRE
CENTRE SCIENTIFIQUE IBM FRANCE
IBM FORSKNINGSSENTERET I BERGEN
WISSENSCHAFTLICHES ZENTRUM DER IBM DEUTSCHLAND
CENTRO DE INVESTIGACION UAM-IBM
BERGEN
WINCHESTER
PARIS
HEIDELBERG
MADRID
PISA
ROME
HAIFA
KUWAIT CITY
CAIRO
TOKYO
مركز الكويت العلمي - ا ب م
CENTRO RICERCA DI PISA
CENTRO RICERCA DI ROMA
המרכז המדעי י.ב.מ. ישראל בע"מ
مركز القاهرة العلمى
G321-0097

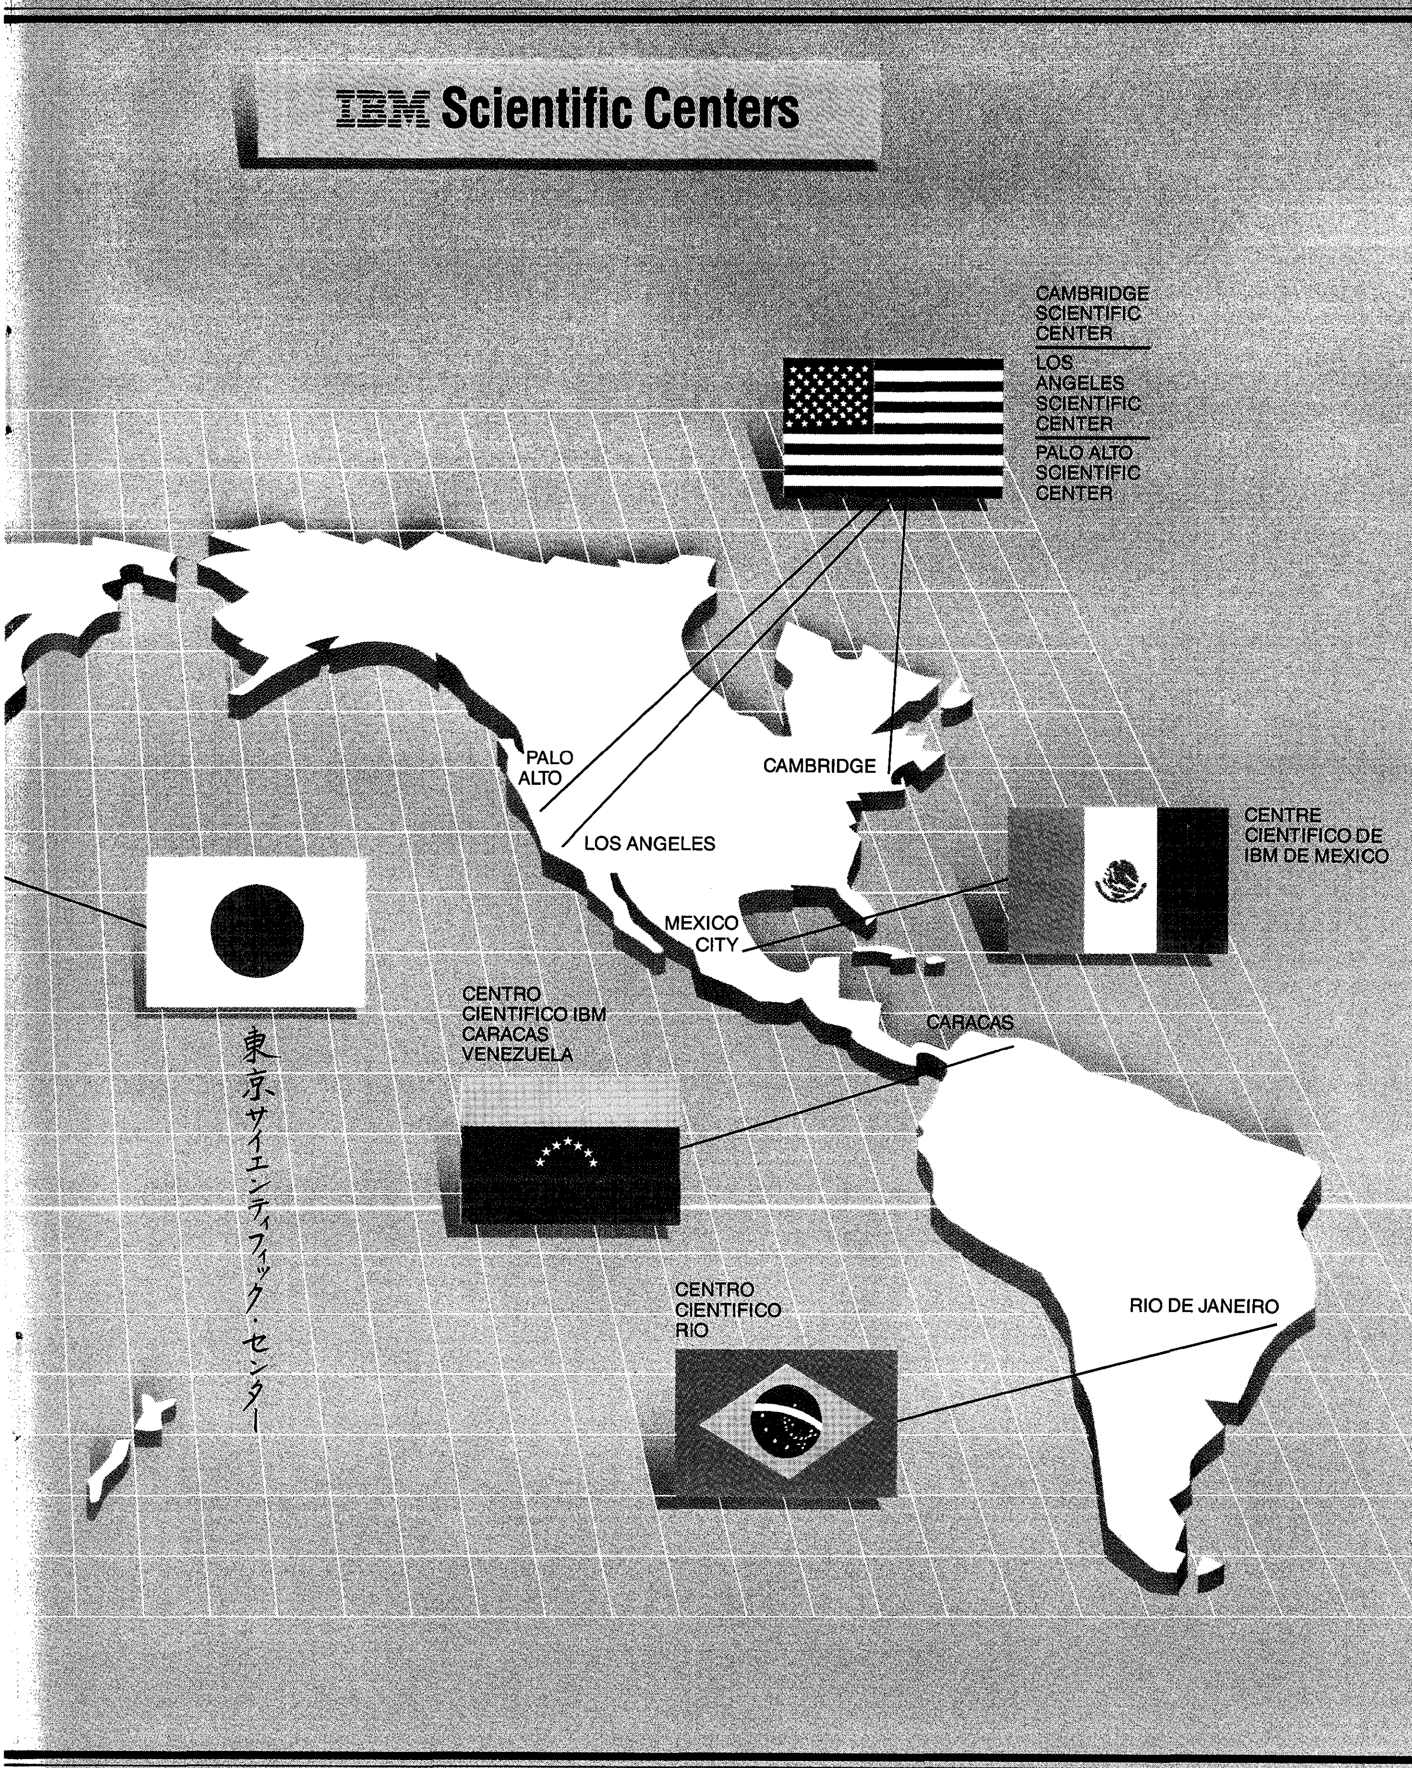
IBM Scientific Centers
CAMBRIDGE SCIENTIFIC CENTER
LOS ANGELES SCIENTIFIC CENTER
PALO ALTO SCIENTIFIC CENTER
PALO ALTO
LOS ANGELES
CAMBRIDGE
CENTRE CIENTIFICO DE IBM DE MEXICO
MEXICO CITY
CENTRO CIENTIFICO IBM CARACAS VENEZUELA
CARACAS
東京サイエンティフィック・センター
CENTRO CIENTIFICO RIO
RIO DE JANEIRO

Figure 1 (A) An original report cover featuring a bubble chamber photograph, emphasizing physics; (B) a report cover from today; (C) announcements of a field-developed program, and (D) an installed user program. The Scientific Centers have generated many such programs.

entific Center in Italy; the Numerically Intensive Computation centers in the United States, Europe, and Japan; and the European Networking Center in West Germany. While they are different organizations, they carry on the same traditions of scientific exploration and service to IBM and its customers.

Contributions

Some further major contributions of the Centers include the virtual machine operating system concept which became Virtual Machine/370 (VM/370), VNET and BITNET, the vector and parallel FORTRAN compilers for the IBM 3090, A Programming Language (APL) and related products, the first IBM Personal Computer (PC) prototype (SCAMP), the Expert Systems Environment (ESE), and the Advanced Interactive Executive/370 (AIX/370) distributed operating system. Two other significant projects are the numerical enhancement of satellite photographs (Figures 2 and 3) and the restoration of famous

Figure 2 LANDSAT image of San Francisco Bay area

Figure 3 Image of volcano dust cloud, Mt.St.Helens eruption, May 18, 1980. Produced from satellite data at the Palo Alto Scientific Center.

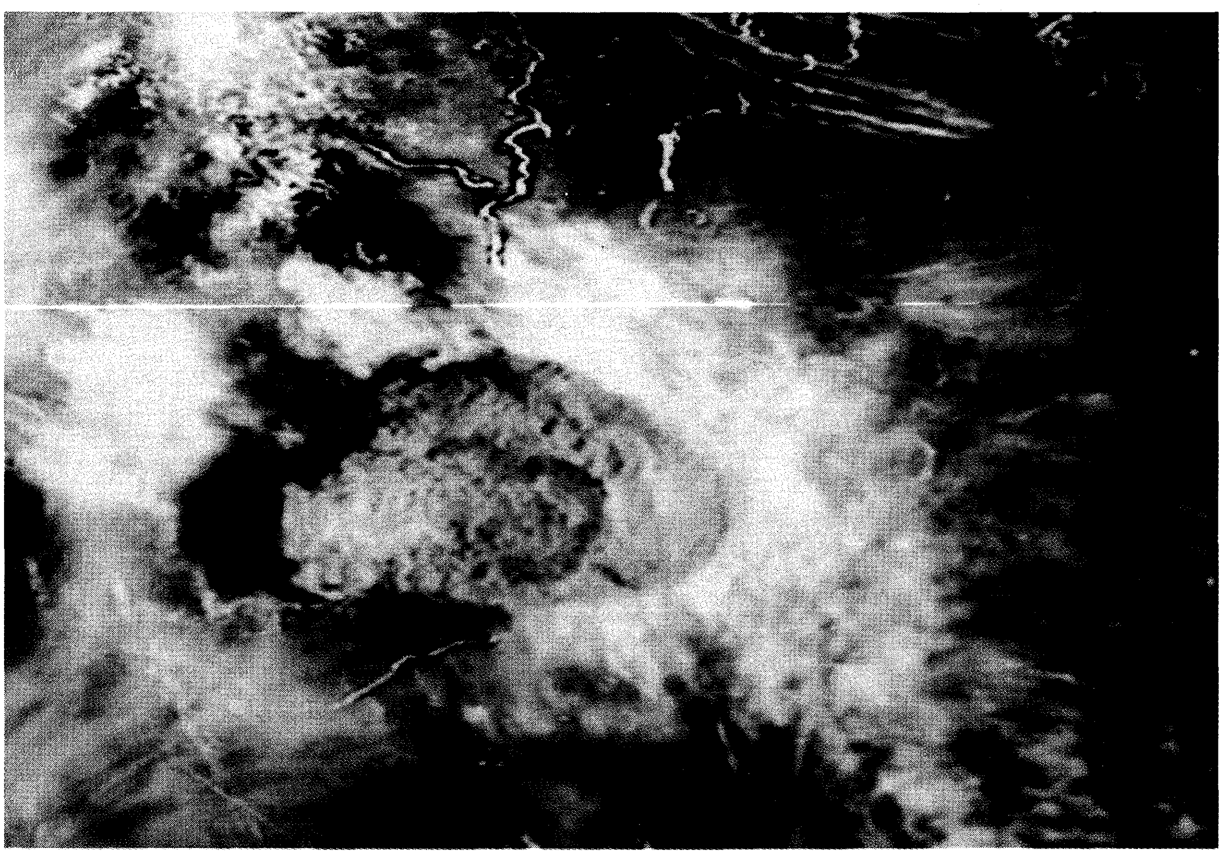

Figure 4 Mona Lisa studied by modern image processing techniques at the Palo Alto Scientific Center.

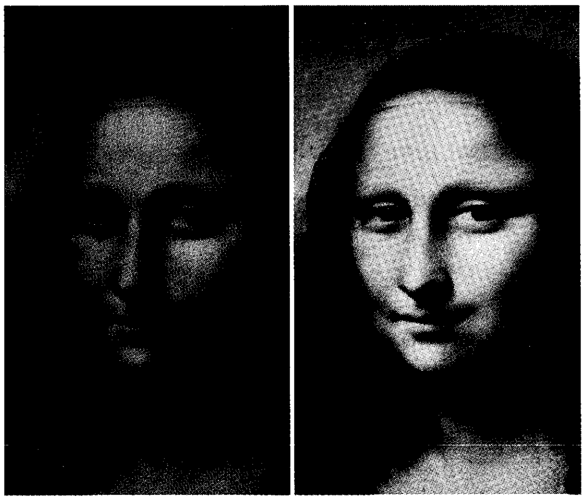

paintings (Figure 4). The Centers have also contributed to solutions for environmental and social problems of their host countries. One example is a computer aid that helps deaf children to speak (Figure 5).

We call special attention to the Hacienda project as an excellent example of cross-center activity and support involving many Centers. The project basis was development, starting in 1978, of an experimental interactive, high-resolution color display and associated software. It led the way for IBM in image processing and resulted in the IBM 7350 Image Processing System, which is still widely used around the world.

Scientific and technical work on the leading edge of change has always been problematical and high risk. The original and continuing mission of the Centers addresses this situation squarely, not only in the choice of problems to be addressed but in an understanding of the special environment within which these problems are best solved. The Centers have been successful in undertaking this work and in transferring technology throughout IBM—frequently into IBM products.

During this first quarter century, there have been over 600 individuals at 26 IBM Scientific Centers worldwide who have been members of the scientific and support staffs. At the same time a greater number of students, temporary assignees, visiting scien-

Figure 5 Deaf child working with a therapist and an IBM Personal Computer. Speech training system developed at the Paris Scientific Center. (From the IBM 1985 Annual Report.)

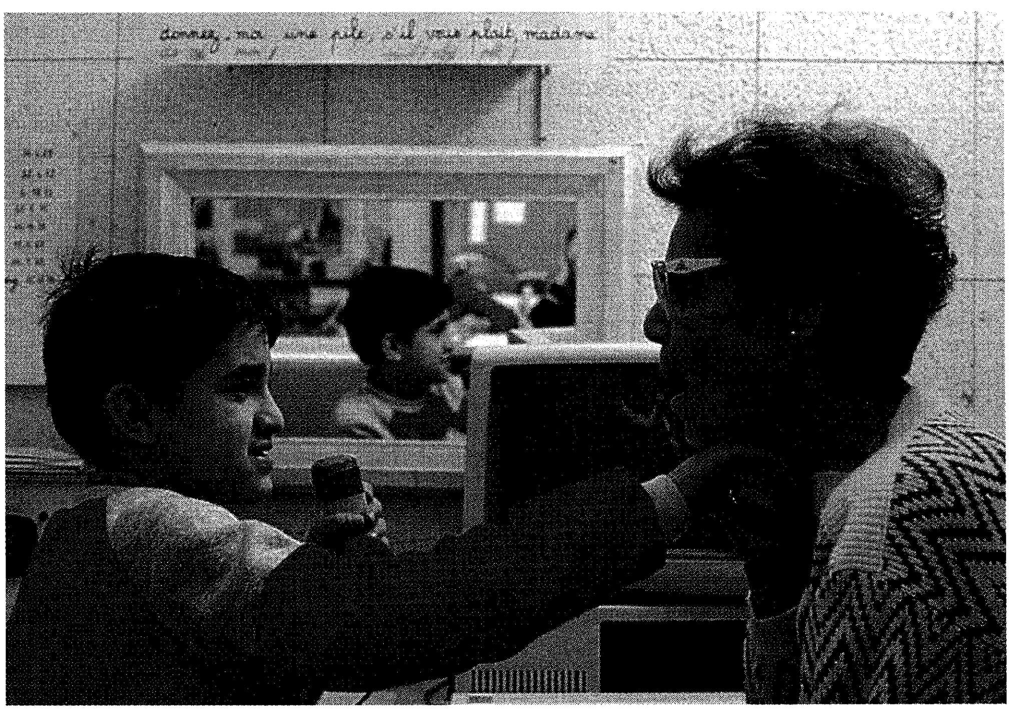

tists, and joint study partners have also participated. It is clearly impossible for us to properly recognize the work of all these dedicated individuals in this paper. The following section gives a brief statement of the major accomplishments and areas of research and development addressed by the dedicated and diverse groups at each Center. Finally, we include a list of selected publications from each of the Centers. The authors acknowledge their considerable debt to all who contributed to these descriptions and particularly to the Center managers who so willingly contributed materials at our request.

The IBM Scientific Centers can be justifiably proud of their record of a quarter century of service to IBM's customers on a wide range of scientific and technical projects. Perhaps more important than the story of the individual Centers has been their common mission, which allowed them relative freedom within IBM to explore new technical areas, take a broad view of research and applications in science and technology, and make significant contributions in areas of growing importance today.

Major accomplishments and projects

Current IBM Scientific Centers

Bergen, Norway (Established 1986)

In 1986 the newest Center was formed, the IBM Bergen Scientific Center. Its focus has been on project areas of importance to Norway, such as information technology, offshore technology, biotechnology, and fish farming. The resulting technical areas under study are vector multiprocessing, parallel processing for the IBM Advanced Interactive Executive™ (AIX™), workstations, mapping, technical documentation, visualization, and animation. In July 1989 the Center obtained the international mission for environmental sciences within IBM. The Center will be the focal point in IBM for environmental modeling.

Cairo, Egypt (Established 1983)

The IBM Cairo Scientific Center was formed through an agreement with the Egyptian Academy of Scientific Research and Technology and began its work in 1983. The initial projects were directed toward environmental studies with various governmental institutions. These projects included modeling the development of Egypt's "new lands," liquid natural gas spills in the Suez Canal, and employment in Egypt. Two new areas emerged over time: Arabic related research such as natural-language processing and signal processing and pattern recognition, as well as aid for the rehabilitation of people with visual and auditory disabilities. Part of the work on natural-language processing is the subject of a paper in this issue.

25 YEARS
Cambridge, Massachusetts (Established 1964)

The IBM Cambridge Scientific Center (CSC) was founded in 1964 as one of the first four Centers. It evolved from a joint project with the Massachusetts Institute of Technology (MIT), which had produced the Compatible Timesharing System (CTSS) for the IBM 7094. From its start and continuing through today, the CSC has focused on machine architecture, interactive computing, programming languages, networking, and system performance. The major contributions of the Center are summarized below.

Virtual Machine/370 (VM/370): The most widely used System/370 operating system, which grew out of the 1968 development of the Control Program 67/Cambridge Monitor System (CP-67/CMS). VM/370 became a product in 1972. After that announcement, further advanced technological contributions were made in the following areas. Multiprocessor support was added, and major performance and scheduling functions were supplied via the Fair-Share Scheduler, both in 1976. That scheduler quickly became the standard VM scheduler. In 1983 the VM Performance Planning Facility (VMPPF) was added for predicting workload under various system configurations. Between 1979 and 1981, VM was enhanced to allow the Multiple Virtual Storage (MVS) operating system to run under VM as a "preferred guest." This work has been carried on in the Processor Resource/Systems Manager™ (PR/SM™) on the IBM 3090.

Remote Spooling Communications Subsystem (RSCS): A peer-to-peer, store-and-forward network protocol, built using virtual machine principles. It is the major network facility used by IBM for internal communications, connecting over 3400 network computers, and is known informally as VNET.

BITNET: A major university computer network, built using RSCS protocols. It has 441 institutions and 1622 nodes in the United States. Its analogues in other countries (EARN, NETNorth, and other academic networks in the Far East and South America) add an estimated 686 institutions and 1036 nodes.

Yale ASCII terminal system: A collaboration with Yale University to permit full duplex ASCII capability for the System/370 through a Series/1 front end, without changes to host programming. The IBM 7171 and then the IBM 4994 replaced the Series/1 and extended this protocol conversion function.

IBM Personal Computer (PC): A series of projects that resulted in Asynchronous Communication Support, IBM 3101 Emulation Support, PC/370 introduced on the PC/XT, PC/IX, and keyboard standards.

Unattended computer operations: An effort to provide remote and unattended distributed operation for midsize computers. The technologies developed appeared in the IBM 4300, 9370, and Application System/400™ (AS/400™).

Operating systems: Early work on UNIX® implementations for IBM systems including the first VM product, Virtual Machine/Interactive eXecutive (VM/IX).

X-Windows: Two projects to make the X-Window System from the Massachusetts Institute of Technology (MIT) available on VM and MVS.

Text processing: A long-standing involvement in text handling has resulted in products such as SCRIPT, which was the precursor for the later Document Composition Facility (DCF), IBM Usability Aid (QPRINT) for the IBM 6670, and Math Formulator for the IBM PC.

Cambridge Control Unit (CCU) and Continuously Executing Channel Interface (CETI): Two efforts that enhanced the ability to connect IBM System/370 computers directly onto high-speed local area networks.

POLITE: An early effort on WYSIWYG ("What you see is what you get") text systems for the System/370 and IBM PC, which influenced the IBM DisplayWrite/370. This text product pioneered the concept of "UNDO" and "REDO" for backtracking.

Caracas, Venezuela (Established 1983)

In 1983 the IBM Caracas Scientific Center opened its doors, with a focus on remote sensing and image processing. Now the emphasis is on three areas of scientific study: astrophysics and computational chemistry, algorithms and data structures, and software engineering. The first has resulted in new information about electron impact excitation and radiative atomic data. The second has seen improved time and space efficiency for generalized binary search trees, the transforming of multiway trees into practical external data structures, and new methods for collision resolution in hashing algorithms. The third has added a new *in situ* distributive sorting algorithm—EXTQUICK—to the tools available to the software engineer, and currently focuses on programming environments.

Haifa, Israel (Established 1972)

Projects of special interest to the planning, medical, and agricultural concerns of Israel have been the major theme in the work of the IBM Haifa Scientific Center, established in 1972. The positive effect of this work is reflected in the progress cited below. Present advanced technology work is focused on image processing for archival applications, scheduling and routing problems, and natural-language processing in Hebrew.

Regional planning: Efforts have contributed to regional rural development through a computer-aided modeling and planning system based on mathematical models and linear optimization. This work has been used by Rehovot, Israel (for whom it was originally developed), Spain, and Costa Rica.

Medical applications: Projects have made noticeable progress on improved diagnostic capabilities for ultrasound testing using signal and image processing techniques, estimating cardiac output and gas content of blood, specialized interactive systems for diagnosis of endocrine disorders, and the further use of signal processing to appreciably decrease background noise (the "cocktail party" noise problem) in assisting the hearing impaired.

Agriculture: Contributions are most notable in the development of an interactive system for constructing irrigation time tables. This system is now used throughout the country and has an important role in the efficiency of agricultural production in their arid lands.

Mathematics: Theoretical and practical applications of numerical analysis to problems arising in mathematical physics and engineering, such as stiff differential equations and partial differential equations with complex boundary conditions.

Heidelberg, Federal Republic of Germany (Established 1968)

When the IBM Heidelberg Scientific Center (HDSC) first opened in 1968, its projects concerned the physics of bubble chambers and scintigraphic photo analysis. Now the areas of interest are more closely aligned with data processing: natural-language capabilities, knowledge-based querying, relational data models, visual languages for end users, additions to A Programming Language (APL) products, numerically intensive computing (NIC), and immunological information processing for transplant surgery. Each of these subjects is detailed below. HDSC is also part of the European Research Initiative EUREKA.

Natural languages: An early project to permit limited natural-language interaction for database queries, with support for many western languages.

LEX: A project within the Institute for Knowledge Based Systems, that is part of HDSC. It is a German language query system for information on the law.

Advanced Information Management Prototype (AIM-P): A project to add an extended Non First Normal Form (NF^2) to the relational data model in IBM's Structured Query Language (SQL).

Extensions to Query-By-Example (QBE): A longstanding developmental effort to add visual programming for end users to relational database systems. More recently, human factors investigations have led to development of a methodology for solving application problems through graphical manipulation.

APL extensions: As a result of first-hand, concentrated use of APL at the HDSC, additional capabilities were developed for APL and APL2 in support of complementary functions, interlanguage communication, and graphic presentation.

NIC applications: Projects with universities and commercial customers to use NIC for reactive flow; chemical, mechanical, and financial modeling; and simulation, in concert with the European Supercomputing Initiative.

Immunology: An effort to support immunological research in transplant surgery across 300 transplant centers worldwide, using an interactive and distributed information system.

Kuwait City, Kuwait (Established 1980)

The IBM Kuwait Scientific Center operates in cooperation with the Kuwait Institute for Scientific Research (KISR) and has been located on its grounds since opening in 1980. Areas of study at the Center are the environment, Arabic languages, and image processing, as presented below.

Environmental studies: Work has been completed on passive cooling and air pollution, and the focus is now on oil pollution and source identification. Another effort is focused on desalinization by reverse osmosis.

Arabic language processing: Projects are under way in IBM Personal Computer-based speech synthesis, speech and printed character recognition, a recent comprehensive Arabic morphological analyzer and generator, course authoring, and desktop publishing with advanced Arabic printing.

Image processing: Ongoing efforts to map Kuwait using LANDSAT 5 satellite data, to study oil distillation columns using thermal infrared imaging, to model pores in catalysts, and in numerical geophysics.

25 YEARS
Los Angeles, California (Established 1964)

The IBM Los Angeles Scientific Center (LASC) is one of the first four Centers founded in 1964. It had its beginning in the joint IBM and University of California at Los Angeles (UCLA) Western Data Processing Center (WDPC), a primary education facility for academic computing as well as for IBM. In the early years, projects emphasized management science, education, engineering, and science. In recent years, some of the projects still center on those areas, and new areas have been added, as described below.

Direct numerical control: Projects on remote computer control of machine tools, with software to cut metal according to customer designs.

Computer-aided instruction (CAI): A number of projects in the development and use of CAI were sponsored, including the development of an interactive supervisor for the System/360.

Oceanography: A joint project with the Scripps Institution to resolve problems in the processing of oceanographic data.

Geophysics: Theoretical studies of earthquake structure using computer simulation and complex models of the earth.

Natural-language processing: Initially a study of applications in medical records, using new syntactic and semantic methods for natural-language information storage and retrieval. Currently, text selection and understanding are being investigated, along with machine translation among many natural languages.

Systems and end users: Projects in operations research (such as a distribution system simulator product), computer-aided design and computer-aided manufacturing (CAD/CAM), application development environments, database security, and distributed computing.

Multidimensional graphics: Joint studies with several universities on the perception of geometric objects in higher dimensions led to advances in computational geometry, displays for exploratory statistical analysis, and algorithms for air traffic and robotics control.

Engineering and scientific systems: Exploration of workstations for use by meteorologists, chemical engineers, and biologists, and for general engineering. Studies in large-scale scientific computing architecture and applications, and in a prototype parallel data store.

Robotics: Projects in vision systems for electronic parts identification and location. Studies are being conducted on artificial intelligence techniques for computer-integrated manufacturing (CIM).

Enterprise and strategic information management: Efforts in improving the effectiveness of business systems planning and integrating it more fully into corporate planning, budgeting, and decision making. A worldwide network of universities and corporations now uses some of the results.

A prototype expert system, called Business Strategy Advisor (BSA), was used by an IBM laboratory to assess competitive software products. Business modeling and a prototype expert system to help identify opportunities for information technologies in enterprises comprise the rest of the work in this area.

Visual programming: Projects to represent programs as forms, with filling-in forms as the way in which a user would see programming. This work is the subject of a paper in this issue.

Security: Studies of tools to enhance the security of IBM's internal computer network and of expert systems for computer security. A prototype system for security audits is now being tested.

Madrid, Spain (Established 1972)

The Autonomous University of Madrid is host to the IBM Madrid Scientific Center, which was chartered in 1972. Major projects at the Center are the automation of the Indies Archive, natural-language processing for Spanish and Catalan, image processing in biology, EXPO'92 visitors information system, 1992 Olympic Games commentators system, and A Programming Language (APL) interpreters.

Image and database processing: The automation of the Indies Archive project is the development and installation of an information system for the management of historical records, integrating a textual database and an archive of pictorial information. A project on image processing in biology, using an IBM PS/2® Model 80, has been carried out in cooperation with the Spanish Molecular Biology Center. Its aim has been to develop algorithms for the filtering, three-dimensional reconstruction, and graphic representation of images of viral particles.

Language processing: The MENTOR (Multitarget ENglish TranslatOR) project is in support of translating IBM manuals from English to Spanish. The approach is valid for other languages, since it relies mainly on a declarative bilingual dictionary that stores most of the information relevant for translation. A lexicon for spelling verification of Spanish and a synonym dictionary have been implemented, together with a morphological analyzer and generator.

EXPO'92 and 1992 Olympic Games information systems: The development of a central information system to provide EXPO'92 visitors with relevant information regarding the fair, events, and activities that will take place in Seville. A 1992 Olympic Games commentators system will provide real time information related to the sports events and access to the Olympic Games Information Systems database to ratio and television sports commentators, using PS/2s.

APL interpreters: Extensive projects, in cooperation with IBM Japan, that have resulted in several releases of APL and APL2 interpreters for PC APL and APL2, and for IBM 5550 and JX Nihongo (Japanese) APL.

Hacienda (IBM 7350): One of the Center partners in the production of the High Level Image Processing System (HLIPS) for Hacienda.

Mexico City, Mexico (Established 1971)

In 1971 this Center was called the IBM Latin American Scientific Center; in 1973 the name was changed to the IBM Mexico Scientific Center, in keeping with its close ties to Mexican projects. Major project areas include artificial intelligence and expert systems, image processing and remote sensing, geographical information systems, statistics and applied mathematics, and database design. Of special interest is the work on geographical information systems that support decision making concerning natural resources by allowing the mixing of different kinds of mapped data, such as temperature, rainfall, and coastal dune vegetation. A number of programs have also been built for the IBM Personal Computer and Personal System/2® (PS/2) for time series analysis, satellite data processing, and database consistency checking.

25 YEARS

Palo Alto, California (Established 1964)

Founded in 1964 as one of the first four Centers, the IBM Palo Alto Scientific Center (PASC) began its history with three major projects: laboratory automation, nuclear power research, and applied physics in materials science and atmospheric physics. Work has also focused on microcode, image processing, and FORTRAN. Other projects came from Washington in 1969 and Houston in 1974 when those Centers closed. The projects that have been significant through the years are listed below.

Large-scale computing: Areas of study have been high-energy, reactor, atmospheric, and plasma physics; geophysics; and aeronautical, chemical, and power engineering. Commercial data processing problems were also addressed. Notable contributions were the Fast Fourier Transform algorithms for solving partial differential equations, and a number of innovative algorithms for vector and parallel machine architectures. Research has ranged from basic physical processes, through model formulation, simulation, numerical analysis, programming, and graphical display techniques to the verification of models by means of measured data.

FORTRAN: Since 1976, when a project to produce a new optimizer was initiated, FORTRAN projects have provided important contributions to performance, data striping, the influential vectorizing FORTRAN compiler for the IBM 3090 Vector Processor, and, most recently, a parallel FORTRAN compiler prototype for use with multiple IBM 3090s.

Microprogramming: Successful work on implementing large system functions using writable control store was first shown within IBM at PASC. Further projects led to the high-level language machine for A Programming Language (APL), microcode performance assists for Virtual Machine Facility/370 (VM/370) and its Conversational Monitor System (CMS), and APL microcode assists for specific models of the System/370, which became the design for the VSAPL product.

IBM Personal Computer (IBM PC) prototype: An outgrowth of work on microprogramming was SCAMP (Scientific Center APL Machine Prototype), the prototype for the IBM 5100 portable desktop computer and a model for later IBM PCs. SCAMP is now in the collection of the Smithsonian Institution.

Expert Systems Environment (ESE): A development effort that added ease-of-use and new implementation techniques to products based technically on Stanford University's EMYCIN.

Advanced Interactive Executive/370™ (AIX/370™): Efforts that built on earlier university work in distributed operating systems and resulted in the introduction of that technology to IBM.

Image processing: Since 1979 a series of image-based projects with the Hacienda (IBM 7350) and joint studies have resulted in applications such as the DIMAPS imaging system, processing of LANDSAT images, an imaging system for the IBM Personal Computer, computer analysis of the Mona Lisa, and medical imaging.

Laboratory automation: In conjunction with Stanford University and the Stanford Linear Accelerator

Center (SLAC), computers monitored high-energy physics experiments and provided immediate graphic output of the results.

Graphics Program Generator: A development effort that resulted in four products and defined new application areas in geographic information systems for public utilities, manufacturing, and process industries.

Networks and software: Programming projects involved the rework of the APL Departmental Reporting System and improvements in the note handling and cross-system communications of PROFS™. Early work in local area networking led to IBM's first Ethernet™ local area network, contributed significantly to technology used in PCnet, and allowed widely dispersed people to communicate inexpensively with the PSInet computer conferencing system.

Paris, France (Established 1977)

The IBM Paris Scientific Center was chartered in 1977, but it built on a tradition of projects that began in 1960 with the IBM France Scientific Development Department. The main areas of exploration since the Center was founded have been mathematics, linguistics, image processing, and artificial intelligence, as described here.

Mathematics: Early work focused on qualitative statistics for analysis of qualitative variables. More recent projects are on precision in complex and lengthy computations, and in understanding chaos.

Linguistics: The deaf children project resulted in the IBM SpeechViewer™ product in 1988. The stenotypy project produced TASF (*Traduction Automatique de Stenotypes en Francais*), a product for creating natural text out of stenotypy, in French. The French thesaurus project developed lexicographic techniques for multiple native languages, that in turn became part of the IBM DW3700 advanced text processing products. The latest work is on automatic dictation capabilities.

Hacienda (IBM 7350): The first Center partner to receive the hardware prototype and the Basic User Software (HBUS), in 1981. The projects at Paris resulted in the High Level Image Processing System (HLIPS) and APL Image Processing Attachment Support (APLIAS), which both became products.

Artificial intelligence: Very early work resulted in the pioneering expert system *Bateau sans medecin* (ship without a medic). More recent work focuses on theoretical breakthroughs in nonclassical logic through development of a qualitative logic model allowing order-of-magnitude reasoning, and resulting systems for nonshallow reasoning.

Programming in Logic (PROLOG) for VM and MVS: Developed in Paris for using logic programming on IBM machines.

Pisa, Italy (Established 1971)

In 1965 the Centro Studi IBM was established in Pisa and became the basis for the IBM Pisa Scientific Center in 1971. This Center has had major project activity in surface hydrology, computer networks, image processing, and econometric software. A present major concern is with language technology. Projects in these and other areas are described below.

Surface hydrology: Project work was stimulated by serious flooding of the Arno River in 1966 and resulted in a model of the river basin that simulates rainfall and runoff. This work was also used for study of the Nile River and Lake Aswan, in Egypt.

RPCNET: An effort to develop a network for VM-based computing centers. This successful network is still running today.

Econometrics: The design and implementation of a complete set of programs to estimate, validate, analyze, and simulate both linear and nonlinear macroeconometric models.

Hacienda (IBM 7350): Design and implementation of the Host Basic User Subroutines (HBUS), which is the host-resident subroutine library for the IBM 7350 Image Processing System.

Medical imaging: Analysis of radiographs of the internal innervation and blood circulation of the cardiac muscle, with experimental work on two- and three-dimensional object detection and recognition.

Volcanic risk modeling: A project for automatic drawing of volcanic eruption risk maps, which were validated using historic eruptions of Mounts Vesuvius, Etna, and St. Helens.

Language technology: Sistema C product development that performs sophisticated interactive checking and correction of Italian texts. Sistema L prototype for printed document reading by scanner, helped by ad hoc lexical tools: it represents a reference point in the area of the OCR-based reading tools for performance in terms of correctness and usability. Further work, in its early stages, aims at human-aided machine translation of English into Italian.

Rio de Janeiro, Brazil (Established 1986)

The IBM Rio Scientific Center was established in 1986 when the Center in Brasilia moved to Rio de Janeiro and integrated all research and educational activities of the Latin American Systems Research Institute, the Software Technology Center, and the IBM Brazil Sumare Plant. Priorities are now more business oriented, and current research areas are databases, data communications, artificial intelligence, logic programming, microelectronics, software engineering, vector processing, and advanced signal processing.

Significant contributions have been made through development of an accurate model of the borer insect and sugarcane ecosystem for plague control, and through study by remote sensing of the environmental impact of the Tucurui hydroelectric facility. Contributions to computer science include the PACCHIP system, an integrated computer-aided design (CAD) package for very large-scale integrated (VLSI) chip design, the CHRIS software tool for database conceptual design and prototyping, an expert system for heuristic learning of medical diagnoses, data compression algorithms for digital images storage and transmission, and halftoning algorithms for laser printers.

Rome, Italy (Established 1979)

The IBM Rome Scientific Center was chartered in 1979 and absorbed the projects and staffs of the Centers at Bari and Venice. Over time the project focus shifted from that of the former Centers to new areas, as described here.

Distributed processing: The Virtual Machine Distributed Facility (VMDF) was used on XCF (eXtended Communication Facility) to create a single-image VM system that allowed file sharing between remote virtual machines and catalog sharing between VM control programs. It was operational on the network for the Centers in Italy until 1984.

Hacienda (IBM 7350): Implemented the IBM 7350 Image Processing System and, since 1981, has used its special color capabilities, along with the IBM 3838 Array Processor, in numerical modeling projects such as seismic wave migration, fluid dynamics, and meteorology.

Natural-language processing: Among projects with relevance to Italy, one began in 1984 with the goal of understanding the Italian language and in 1987 resulted in the most complete text understanding system presently available in Italy.

Speech synthesis: A text-to-speech synthesizer project for Italian, begun in 1982, further resulted in speech recognition of 20 000 Italian words by 1988.

ECSEC: Projects in numerically intensive computing (NIC) resulted in the establishment of the European Center for Scientific and Engineering Computation (ECSEC) as a department of the IBM Rome Scientific Center in 1984. Since then, the Center has become the European focal point in IBM for NIC.

Tokyo, Japan (Established 1970)

The IBM Tokyo Scientific Center was founded in 1970. Main project areas have been natural-language support, image processing, graphics and displays, scientific and engineering computation, and education, as described below.

Natural languages: Research conducted in machine translation between English and Japanese. Work has also been done on converting Kana to Kanji for a Japanese word processor.

Image processing: Many joint studies with government and university groups emphasized regional and coastal planning. Recent studies included optical character recognition (printed and handwritten) and color image databases.

Graphics and displays: Designed the application programming interface (API) for the IBM 5080 display and its follow-on with Kingston, developed rendering methods including a new texture mapping, and applied them to visual simulation.

Scientific and engineering computation and visualization: Projects included chemical computer-aided design, computational chemistry, and engineering simulation, such as simulating the air flow in a clean room. New graphics techniques have been applied to the visualization of molecules and air flow.

Education: Having always conducted projects with university partners, recent efforts have involved assistance for the handicapped, mathematical computer-aided instruction, and campus networking.

Winchester, United Kingdom (Established 1979)

In 1979 the IBM United Kingdom Scientific Centre (UKSC) adopted a new charter and a new project focus, and moved from Peterlee to Winchester, England. It remains the only Center with a formal connection to an IBM laboratory—IBM United Kingdom Laboratories Limited, at Hursley Park, England. It also is unusual in its extensive use of visiting professionals, forming roughly half of the staff. The new charter established a new set of project areas for the Center: image processing, graphics systems, and speech processing. The major contributions in these areas are outlined below. In 1988 the Center was a joint winner of the British Computer Society Award for applications of computing.

Interactive Applications eXecutive (IAX): An application of high content, high-resolution image processing capabilities to medical imaging of the heart and neurons, for example.

Graphics: A number of projects surrounding the use of graphical techniques for the presentation of forms and the visualization of large amounts of complex data. Areas for study have been solid geometry, raster display of complex three-dimensional constructions, molecular modeling for the pharmaceutical industry, high-speed imaging using arrays of transputers, automatic interpretation of images, stereoscopic computer vision, fluid flow, astronomy, archaeology, theoretical physics, molecular beam epitaxy, liquid crystal structure, and computer-generated art. The WINchester SOlid Modeling system (WINSOM) and its solid geometry applications are discussed in three other papers in this issue.

Speech processing: Projects focused on English speech synthesis from text and especially automatic phonetic construction from text.

Former IBM Scientific Centers

Bari, Italy (1969–1979)

One of the smallest Centers, the IBM Bari Scientific Center studied the fields of computer-aided instruction and natural-language processing and later relational databases and query languages. It saw the development of the A Query Language (AQL) interpreter as a relational database extension to the A Programming Language (APL) systems. It was consolidated into the new IBM Rome Scientific Center in 1979.

Brasilia, Brazil (1980–1986)

The IBM Brasilia Scientific Center was established in 1980 in the new capital city of Brazil. It's focus was on enhancing the quality of life, the work environment, and the overall community's well-being. In 1984 the Center was honored for this work by receiving the American Chamber of Commerce Award for Corporate Service to the Community. In 1986 the Center merged with other activities in Brazil and moved to Rio de Janeiro.

Grenoble, France (1967–1973)

The first Center outside the United States was the IBM Grenoble Scientific Center. Its efforts were on operating systems and compilers. It played a significant role in the development of CP-67/CMS, in close cooperation with the IBM Cambridge Scientific Center. There was also a project on incremental compilation, the results of which are in use in the FORTRAN Interactive Debug and PL/I Incremental Compiler products. The Center closed in 1973 and much of its work was reassigned to what later became the IBM Paris Scientific Center.

Houston, Texas (1966–1974)

Much of the work at the IBM Houston Scientific Center focused on the use of array processors, such as the IBM 2938 and 3838. These studies included numerical precision, seismic applications, and development of array algorithms. Specific efforts included the array processing language VECTRAN, optical and holographic image processing and display, medical data processing for the heart and X-rays, chemical engineering real-time process control, and electrical power systems analysis. This Center was merged into the IBM Palo Alto Scientific Center in 1974.

New York City, New York (1964–1972)

Founded in 1964, the IBM New York Scientific Center was one of the first four Centers chartered. It grew out of the prototypical Mathematics and Application Department, originally formed in 1955 in New York City. There were two basic areas of concentration: management science, which included studies of simulation, network analysis, and computational programming; and engineering design, which devised nonstandard data acquisition attachment capabilities for the System/360, graphics facilities for use in textile design, and graphical methods for urban planning. This Center became part of the foundation for the IBM Philadelphia Scientific Center in 1972.

Peterlee, United Kingdom (1969–1979)

This Center was established in 1969 as the IBM United Kingdom Scientific Centre. The major project areas were relational databases and local and regional planning. The relational database work culminated in the early and successful "Peterlee Model" that did much to popularize the concept. The planning work involved a number of projects with governments on using computers for planning. In 1979 the Center altered its focus considerably and moved to Winchester.

Philadelphia, Pennsylvania (1972–1974)

In 1972 the IBM New York Scientific Center and the IBM Research Division group on A Programming Language (APL) joined forces and started the IBM Philadelphia Scientific Center. APL became the major project area for the Center, along with the project areas from New York: simulation, network analysis, and computational programming. The work on APL resulted in the APLNET network, an APL business planning system, a long-standing APL computer center and time-sharing service for other IBM groups, extensions to APL itself, and the groundwork in generalized array theory that became the basis for A Programming Language 2 (APL2). Other projects included airline crew and telephone operator scheduling, cash and asset management, earth resources evaluation, and life insurance budget planning. Most of the Center's work was relocated to the Palo Alto APL development group in 1974, and this Center was closed.

Venice, Italy (1969–1979)

The IBM Venice Scientific Center was located in a thirteenth-century palace—the only Center so honored. A storm surge in 1966, and subsequent efforts to save Venice from the sea, resulted in the creation of this Center. As a result, much of its work centered on hydrology and air pollution. The numerical modeling of combined tidal and meteorological effects in the Venice lagoon, and of the sinking phenomenon in the Venice area due to the pumping of groundwater, remain the major contributions of this Center to basic knowledge in these fields. In 1979 the Center became part of the new IBM Rome Scientific Center.

Wheaton, Maryland (1967–1969)

A Center was established in 1967 in Wheaton, Maryland, to be known as the IBM Washington Scientific Center. Significant areas for this Center were environmental sciences, text processing, cartography, and microcode. The environmental sciences area involved work on mathematical techniques for descriptions and mapping of nature, such as weather prediction and water resources—reservoirs, pollution, and waste treatment automation. The Center's work moved to Palo Alto in 1969.

Selected publications

Current IBM Scientific Centers

Bergen, Norway

L. de Pillis Lindheim, J. Petersen, and J. de Pillis, "An iterative solution to special linear systems on a vector hypercube," *3rd Conference on Hypercube Concurrent Computer and Applications, Vol. II*, ACM Press (1988).

A. D. Jenkins and J. H. Salvesen, "Ocean current modelling at the Bergen Scientific Centre," *ODAP Current Modelling Workshop, Roros, 1988*, Trondheim: Oceanor—MOMOP Report No. 1 (1988).

A. Kamel, "Elastic modeling on the IBM 3090 vector multiprocessor," *European Conference on Mathematics in Industry*, Glasgow (August 28–31, 1988).

Cairo, Egypt

K. Abdel Hamid, M. Hashish, O. Emam, D. Kamal, A. El-Hetw, and K. Ahmed, "A multi-feature computer-based speech-training system for deaf children," *IEEE Engineering in Medicine and Biology (EMBS) 9th Annual Conference*, Boston (November 1987).

M. Elghonemy, M. Fikri, M. A. Hashish, and E. Talkan "An iterative method for formant extraction using zero crossing interval histograms," *IEEE Melcon '85*, Madrid, Spain (October 8–10, 1985).

Hisham El-Shishiny and Bayoumi Attia, "Multi-objective modeling for the planning and management of new lands in Egypt—A case study," *IFAC Conference on Systems Analysis Applied to Water and Related Land Resources*, Lisbon (October 1985).

O. S. Emam and M. A. Hashish, "Application of hidden Markov models to the recognition of isolated Arabic words," *10th National Computer Conference and Exhibition*, Jeddah (February 28 to March 2, 1986).

M. A. Serag El Din, "Spread of NG following a small LNG spill from a ship-tanker," *Proceedings of the 4th International Conference on Numerical Methods in Thermal Problems*, Swansea (July 1985).

Cambridge, Massachusetts

Y. Bard, "The VM Performance Planning Facility (VMPPF)," *Computer Measurement Group (CMG) Transactions* **53**, 53–59 (Summer 1986).

R. J. Creasy, "The origin of the VM/370 time-sharing system," *IBM Journal of Research and Development* **25**, No. 5, 483–490 (September 1981).

E. C. Hendricks and T. C. Hartmann, "Evolution of a virtual machine subsystem," *IBM Systems Journal* **18**, No. 1, 111–142 (1979).

L. H. Holley, R. P. Parmelee, C. A. Salisbury, and D. N. Saul, "VM/370 asymmetric multiprocessing," *IBM Systems Journal* **18**, No. 1, 47–70 (1979).

F. T. Kozuh, D. L. Livingston, and T. C. Spillman, "System/370 capability in a desktop computer," *IBM Systems Journal* **23**, No. 3, 245–254 (1984).

R. A. Meyer and L. H. Seawright, "A virtual machine time-sharing system," *IBM Systems Journal* **9**, No. 3, 199–218 (1970).

Daniel J. Oberst and Sheldon B. Smith, "BITNET: Past, present, and future," *EDUCOM Bulletin* **21**, No. 2 (Summer 1986).

R. P. Parmelee, T. L. Peterson, C. C. Tillman, and D. J. Hatfield, "Virtual storage and virtual machine concepts," *IBM Systems Journal* **11**, No. 2, 99–130 (1972).

L. H. Seawright and R. A. MacKinnon, "VM/370—a study of multiplicity and usefulness," *IBM Systems Journal* **18**, No. 1, 4–17 (1979).

Caracas, Venezuela

W. Cunto and J. L. Gascon, "Improving time and space efficiency in generalized binary search trees," *Acta Informatica* **24**, 583 (1987).

W. Cunto, G. H. Gonnet, and J. I. Munro, "EXTQUICK: an *in situ* distributive external sorting algorithm," *Proceedings VI International Conference of the Chilean Computer Science Society* (1986), p. 61.

W. Cunto and P. V. Poblete, "Two hybrid methods for collision resolution in open addressing hashing," *Scandinavian Workshop on Algorithm Theory*, Lecture Notes in Computer Science 318, Springer-Verlag, NY (1988), p. 113.

C. Mendoza, "Electron impact excitation of 2p q and 3p ions," *Atomic Data Workshop: Low Energy Collision Theory Techniques for Atomic Excitation and Radiative Data*, W. B. Eissner, Editor, Proceedings of the DL/SCI/R24 (SERC Daresbury Laboratory) (1986), p. 80.

C. Mendoza and C. J. Zeippen, "Radiative atomic data for neutral magnesium. I. Oscillator strengths," *Astronomy and Astrophysics* **179**, 339 (1987).

Haifa, Israel

A. Angel, S. Gal, and J. Raviv, "Optimal procedures for sequential discrimination between two classes," *Proceedings of MEDINFO 80*, Tokyo (October 1980).

Y. Ben-Ari and S. Gal, "Optimization of replacement policy in a dairy herd," *European Journal of Operations Research* **23**, No. 2 (February 1986).

D. Chazan and S. Gal, "A markovian model for a central blood bank," *Management Science* **23** (January 1977).

D. Chazan, Y. Medan, and U. Shvadron, "Noise cancellation for hearing aids," *Proceedings of ICASSP 86*, Tokyo (April 1986).

I. Efrat and M. Tismenetsky, "Parallel iterative linear solvers for oil reservoir models," *IBM Journal of Research and Development* **30**, No. 2, 184–192 (March 1986).

S. Gal, M. Landsberger, and B. Levykson, "A compound strategy for search in the labor market," *International Economic Review* **22**, No. 3, 597–609 (1981).

E. Gofman and M. Rodeh, "Incorporating unknown pipe characteristics into loop oriented hydraulic network solvers," *Journal of the Hydraulics Division, ASCE*, No. HY9, 107 (September 1981).

M. C. Golumbic, M. Markovich, U. J. Schild, and S. Tsur, "A knowledge-based expert system for student advising," *IEEE Transactions on Education* **E-28**, 120–124 (March 1986).

D. Ramm and D. Chazan, "A mixed numerical-analytical method for underground water flow and tracer displacement," *Water Resources Research* **16** (1980).

E. Walach, J. Bruck, D. Chevion, E. Karnin, and D. Ramm, "A modified block truncation coding technique for image compression," *Proceedings of the International Conference on Advances in Image Processing and Pattern Recognition*, Pisa, Italy (December 1985).

A. Ziv and V. Amdursky, "On the numerical solution of stiff linear systems of the oscillatory type," *SIAM Journal of Applied Mathematics* **33,** No. 4 (December 1977).

Heidelberg, Federal Republic of Germany

B. Alschwee and S. Grundmann, "System design for a computer aided juridical expert system," *Automated Analysis of Legal Texts, Logic, Informatics, Law*, A. A. Martino and F. Socci Natali, Editors, North Holland, Elsevier Science Publishers, Amsterdam, 1986.

U. Arend and J. Wandmacher, "On the generality of logical recording in spatial inference tasks," *Acta Psychologica* **65,** 193–210 (1987).

A. Blaser et al., "Integrated Data Analysis and Management System (IDAMS)—Feature description," in *Relational Data Base Systems*, J. W. Schmidt and M. L. Brodie, Editors, Springer-Verlag, NY (1983).

T. Bollinger, U. Hedtstuck, C.-R. Rollinger, and R. Studer, "Text understanding in LILOG," in *Linguistic Approaches in Artificial Intelligence, Duisburg Papers on Research in Language and Culture* **6,** A. Kunz and U. Schmitz, Editors, Verlag Peter Lang, Frankfurt (1988).

G. Bouma, E. König, and H. Uszkoreit, "A flexible graph-unification formalism and its application to natural-language processing," *IBM Journal of Research and Development* **32,** No. 2, 170–184 (March 1988).

P. Dadam, R. Erbe, J. Guenauer, V. Lum, P. Pistor, G. Walch, H.-D. Werner, and J. Woodfill, "Designing DBMS support for the temporal dimension," *Proceedings of the 1984 ACM SIGMOD Conference*, June 18–21, 1984, Boston, ACM Order No. 472840.

G. Dueck and T. Scheuer, "Threshold accepting—A general purpose optimizing algorithm appearing superior to simulated annealing," to appear in the *SIAM Journal on Computing* (1989).

F. Guenthner, H. Lehmann, and W. Schönfeld, "A theory for the representation of knowledge," *IBM Journal of Research and Development* **30,** No. 1, 39–56 (January 1986).

G. Jaeschke and H.-J. Schek, "Remarks on the algebra of non first normal form relations," *Proceedings of the ACM SIGACT-SIGMOD Conference on Principles of Database Systems*, Los Angeles (March 29–31, 1982).

E. Keppel et al., "Implementation of a data dictionary system to maintain a distributed information system for transplant data," *Lecture Notes in Medical Informatics* **35,** Springer-Verlag, Berlin (1988).

H. Lehmann, N. Ott, and M. Zoeppritz, "User experiments with natural language for data base access," *Proceedings of the 7th International Conference on Computational Linguistics (COLING'78)*, Bergen, Norway (August 14–16, 1978).

H. Lehmann, N. Ott and M. Zoeppritz, "A multilingual interface to databases," *IEEE Database Engineering Bulletin* **8,** No. 3 (September 1985).

V. Linnemann, "Non-first-normal-form relations and recursive queries: An SQL based approach," *Proceedings of the 3rd IEEE International Conference on Data Engineering*, Los Angeles (February 1987).

G. Rohr, "Mental models—On the relation between structure and its symbolic representation in human-computer interaction," *Zeitschrift für Psychologie* **193,** No. 4 (1985).

M. Zoeppritz, "Human factors of a natural language enduser system," *Enduser Systems and Their Human Factors, Lecture Notes in Computer Science*, A. Blaser and M. Zoeppritz, Editors **150,** Springer-Verlag, Berlin (1983).

Kuwait City, Kuwait

O. Alameddine, "Arabesques and Koufi art," *Die Pixel Revolution*, IBM Deutschland, GmbH (1987).

O. Alameddine, M. F. Quinn, K. Al-Jamal, "Experimental measurements of the downward infra red sky radiation in Kuwait," *Applied Optics* **23,** No. 3 (February 1984).

K. Al-Jamal, O. Alameddine, H. Al-Shami, and N. Shaabaan, "Passive cooling evaluation of room pond systems," *Solar and Wind Technology* **5,** No. 1 (1988).

Y. A. El-Imam, "Synthesis of Arabic speech using partial syllables," *Proceedings of the First International Conference on Arabization and Informatics*, AED and IRSIT, Tunisia, Tunis (March 9–11, 1988).

A. H. Kamel, S. E. Akashah, F. A. Leeri, and M. Fahim, "Particle size distribution in oil-water dispersions using image processing," *Computers and Chemical Engineering* **11,** No. 4, 435–439 (1987).

S. M. Metwalli, A. H. Kamel, and A. A. Saheb, "Surface roughness effect on laser speckle spectral density," *SPIE's Technical Symposium Southeast on Optics and Optoelectronic Systems*, Orlando, FL (March 31–April 4, 1986).

M. Quinn and O. Alameddine, et al., "A deconvolution technique for determining the intrinsic fluorescence decay lifetimes of crude oil," *Applied Spectroscopy* **42,** No. 3 (1988).

A. R. Ragab, S. M. Metwalli, and J. Rueda, "Preliminary assessment of image processing application to large deformation measurement," *Proceedings of the Third Cairo University Conference on Mechanical Design and Production*, Pergamon Press, Elmsford, NY (March 1986).

S. Sami, "Towards high quality Arabic script output," *Panel on Non-Latin and Non-alphabetic Scripts, IFIP Congress 86: 10th World Computer Congress*, Dublin, Ireland (September 1–5, 1986).

Los Angeles, California

B. Arbab, "Object identification from parallel light stripes," *Proceedings 10th International Joint Conference on Artificial Intelligence*, 1145–1148 (1987).

R. M. Buckley and W. B. Broadwell, "A dynamic model for contouring NC devices," *The Expanding World of Numerical Control*, M. A. DeVries, Editor, Village Press, NY (1972), pp. 156–161.

R. M. Buckley and J. P. Mayfield, "Applications of a numerical geometry system in engineering," *Proceedings of the IEEE 13th Annual Design Automation Conference* (1976), pp. 23–33.

M. M. Connors, C. Coray, C. J. Cuccaro, W. K. Green, D. W. Low, and H. M. Markowitz, "Distributed system analysis: The distribution system simulator," *Management Science* **18,** B425–453 (1972).

E. H. Coughran, *A Preliminary Report on Shipboard Computing*, Report G320-2615, IBM Los Angeles (1968).

K. Dahlgren, *Naive Semantics for Natural Language Understanding*, Kluwer Academic Publications, Boston (1988).

B. Dimsdale, "A natural language information retrieval system," *Proceedings of the IEEE* **54,** No. 12 (1966).

B. Dimsdale, "On multiconic surfaces," *IBM Journal of Research and Development* **19,** 523–529 (1975).

B. Dimsdale, "Convex cubic splines," *IBM Journal of Research and Development* **22,** 168–178 (1976).

E. B. Fernandez, R. C. Summers, T. Lang, and C. D. Coleman, "Architectural support for system protection and database security," *IEEE Transactions on Computer* **C-27,** No. 8 (1978).

E. B. Fernandez, R. C. Summers, and C. Wood, *Database Security*, Addison-Wesley, Reading (1977).

P. Gongla, G. Sakamoto, A. Back-Hock, P. Goldweic, L. Ramos, R. C. Sprowls, and C.-K. Kim, "S*P*A*R*K: A knowledge-based system for identifying competitive uses of information technology," *IBM Systems Journal* **28,** No. 4, 628–645 (1989); this issue.

A. Inselberg, "Cochlear dynamics: The evolution of mathematical model," *SIAM Review* **20,** 301–904 (1978).

A. Inselberg, "Multi-dimensional graphics," *Proceedings of Eurographics '85*, Nice, France (1985), pp. 315–325.

A. Inselberg, "Intelligent instrumentation and process control," *Proceedings of the IEEE Conference on Artificial Intelligence Applications* (December 1985), pp. 302–307.

A. Inselberg, "Parallel coordinates for multidimensional graphics," *Proceedings of the National Computer Association Annual Conference*, Philadelphia (March 1987), pp. 547–566.

A. Inselberg and B. Dimsdale, "Intelligent manufacturing and integrated instrumentation," *Proceedings of ASME Conference on Integrated Intelligent Manufacturing*, Anaheim, CA (December 1986), pp. 341–357.

S. A. Jurovics, "Solar radiation data, natural lighting and building energy minimization," *Proceedings of ASHRAE* (1979).

D. Knight, M. Harry, J. Howard, and J. Rivero, "Engineering support systems for engineering managers," *IEEE Micro* **5** (1985), pp. 22–26.

P. S. Newman, "Towards an integrated development environment," *IBM Systems Journal* **21,** No. 1, 81–107 (1982).

M. M. Parker, "Enterprise information analysis: Cost-benefit and the data managed system," *IBM Systems Journal* **21,** No. 1, 108–123 (1982).

M. M. Parker and R. J. Benson with H. E. Trainor, "Information economics: Linking business performance to information technology," Prentice Hall, Inc., Englewood Cliffs (1988).

N. C. Shu, "Visual programming," Van Nostrand Reinhold Co., New York (1988).

G. Silverman, R. Tsai, and M. Lavin, "Locating polyhedral objects from edge point data," *Proceedings of the International Joint Conference on Artificial Intelligence* (IJCAI), Milano, Italy (1987).

R. C. Summers, "A resource sharing system for personal computers in a LAN: Concepts, design, and experience," *IEEE Transactions on Software Engineering*, 895–904 (1987).

R. C. Summers, "An overview of computer security," *Computer and Network Security*, M. D. Abrams and H. J. Podel, Editors, *IEEE Computer Society Press* (1987), pp. 9–25.

W. J. Van De Lindt, "Movement of the mohorvicic discontinuity under isostatic conditions," *Journal of Geophysical Research*, No. 4, 72 (1967).

Madrid, Spain

M. Aguilar, A. Garcia, P. J. Pascual, J. Presa, and A. Santisteban, "Computer system for scanning tunneling microscope automation," *Surface Science* **181,** 191–199 (1987).

J. Brown and M. Alfonseca, "Solutions to logic problems in APL2," *ACM SIGAPL*, APL Quote Quad **17–4,** 356–361 (May 1987).

L. E. Donate, L. Herranz, J. P. Secilla, J. M. Carazo, H. Fujisawa, and J. L. Carrascosa, "Bacteriophage T3 connector: three-dimensional structure and comparison with other viral head-tail connecting regions," *Journal of Molecular Biology*, 201 (1988).

L. F. Escudero and M. Rebollo, "A mixed integer programming approach to multispectral image classification," *Pattern Recognition* **9,** 45–57 (1977).

J. Jimenez and M. A. Hernan, "Computer analysis of a high speed film of the plane turbulent mixing layer," *Journal of Fluid Mechanics* **119,** 323–345 (1982).

J. Jimenez, P. Moin, R. Moser, and L. Keefe, "Ejection mechanisms in the sublayer of a turbulent channel," *Physics of Fluids* **31,** 1311–1313 (1988).

J. Jimenez, A. Santisteban, J. M. Carazo, and J. L. Carrascosa, "Computer graphics display method for visualizing three-dimensional biological structures," *Science* **232,** 1113–1115 (1986).

M. Kindelan, "Dynamic modelling of greenhouse environment," *Transactions of the ASAE* **23,** 1233–1239 (1980).

A. Santisteban, L. R. Berriel, and J. Bescos, "Image restoration for a defocused optical system," *Applied Optics* **22,** 2772–2780 (1983).

L. Sopena, "Linguistica y procesamiento de textos en castellano," *PC WORLD Espana* **8,** 63–69 (February 1986).

L. Sopena, "Natural language grammars for an information system," *ACM SIG Information Retrieval* **17,** No. 4, 75–80 (1983).

M. L. Tavera, M. Alfonseca, and J. Rojas, "The IBM Personal Computer APL system," *ACM SIGAPL, APL Quote Quad* **14-4,** 333–337 (June 1984).

M. Udo, Y. Akimoto, S. Kaneko, T. Sanuki, and M. Alfonseca, "Japanese APL language system on IBM multistation 5550," *ACM SIGAPL, APL Quote Quad* **16-4,** 326–334 (July 1986).

Mexico City, Mexico

S. Castillo, E. Gutierrez, and O. Alvarez, "Phenology of coastal dune vegetation," *Proceedings of the X Mexican Botany Congress*, Guadalajara, Mexico (1987).

E. Gonzalez, S. Negrete, and P. Noriega, "An evolving data base design," *Proceedings of the V Meeting of the Mexican AI Association*, Merida, Mexico (1988).

J. Macias and M. Angulo, "A climatic-botanic information system of the state of Veracruz," *Proceedings of the XIII Conference of the International Cartographic Association*, Morelia, Mexico (October 1987).

Palo Alto, California

J. F. Asmus, R. Bernstein, J. V. Dave, and H. J. Myers, "Mona Lisa," *Perspectives in Computing* **7,** No. 1, 11–22 (1987).

J. V. Dave, "Contrast attenuation factors for remote sensing," *IBM Journal of Research and Development* **23,** No. 2, 214–223 (1979).

J. V. Dave, R. Bernstein, and H. G. Kolsky, "Importance of higher-order components to multispectral classification," *IBM Journal of Research and Development* **26,** No. 6, 715–723 (1982).

J. V. Dave and J. Gazdag, "Reduction of random noise from multiband data using phase relationships among their Fourier coefficients," *IBM Journal of Research and Development* **28,** No. 4, 399–411 (1984).

A. A. Dubrulle, *A Version of EISPACK for the IBM 3090VF,* Report G320-3510, IBM Corporation (June 1988); available through IBM branch offices.

A. A. Dubrulle, R. G. Scarborough, and H. G. Kolsky, *How to Write Good Vectorizable FORTRAN,* Report G320-3478, IBM Corporation (September 1985); available through IBM branch offices.

H. P. Flatt, "Computer modeling in energy and the environment," *IBM Journal of Research and Development* **25,** No. 5, 571–580 (1981).

P. J. Friedl, "The use of computer conferencing by a university faculty," IFIP TC6/WG6.5, *Conference on Message Handling Systems* (April 1987), pp. 387–402.

J. Gazdag, G. Radicati, P. Sguazzero, and H. H. Wang, "Seismic migration on the IBM 3090 Vector Facility," *IBM Journal of Research and Development* **30,** No. 2, 173–183 (1986).

J. Gazdag and H.-H. Wang, "Concurrent computing by sequential staging of tasks," *IBM Systems Journal* **28,** No. 4, 646–660 (1989); this issue.

A. Hassitt, J. W. Lageschulte, and L. E. Lyon, "Implementation of a high level language machine," *Communications of the ACM* **16,** No. 4, 199–212 (April 1973).

A. Hassitt and L. E. Lyon, "An APL emulator on System/370," *IBM Systems Journal* **15,** No. 4, 358–378 (1976).

A. H. Karp and R. G. Babb II, "A comparison of twelve parallel FORTRAN dialects," *IEEE Software,* 52–66 (September 1988).

J. Littman, "The first portable computer," *PC World* **1,** No. 7, 294–300 (1983).

L. E. Lyon, "Striped I/O for FORTRAN," Report G320-3513, IBM Corporation (June 1988); available through IBM branch offices.

H. J. Myers, "Compiling optimized code from decision tables," *IBM Journal of Research and Development* **16,** No. 5, 489–503 (September 1972).

R. G. Scarborough and H. G. Kolsky, "Improved optimization of FORTRAN object programs," *IBM Journal of Research and Development* **24,** No. 6, 660–676 (1980).

R. G. Scarborough and H. G. Kolsky, "A vectorizing FORTRAN compiler," *IBM Journal of Research and Development* **30,** No. 2, 163–171 (1986).

IBM Program 5875-EEE, *Personal Computer Image Processing System (PCIPS),* IBM Corporation; available through IBM branch offices.

IBM Programs PRPQ 5799-AYB, *Interactive Geo-Facilities Graphics Support*; 5668-941, *Graphics Program Generator Program Product*; *Graphics One-Line Diagram: Display, Edit, and Query Program (FDP)*; 5668-905, *Composed Document Viewing Utility* (Part of Graphic Data Query Facility, Release 3), IBM Corporation; available through IBM branch offices.

IBM Program 5796-PLN, *A Departmental Reporting System II (ADRS II),* IBM Corporation; available through IBM branch offices.

IBM Program 5798-DRT, *PROFS Note Maintenance Facility,* IBM Corporation; available through IBM branch offices.

IBM Installed User Program 5796-PQW, *SORT/MERGE Performance Improvements,* IBM Corporation; available through IBM branch offices.

IBM Installed User Program 5796-PKR, *FORTRAN H Extended Optimization Enhancement,* IBM Corporation (1978); available through IBM branch offices.

IBM Programs 5799-ALK, IBM PRPQ MF2608 and IBM RPQ 500256, *APL/CMS and APL Assist,* IBM Corporation (1974); available through IBM branch offices.

IBM Installed User Programs 5796-PJB, *APL Decision Table Compiler* (1977); 5796-PKD, *General Cross Assembler Generator* (1978); and 5790-PPJ, *APL Performance Tools,* IBM Corporation; available through IBM branch offices.

Paris, France

M. Abrams, A. Blusson, V. Carrere, T. Nguyen, Y. Rabu, "Image processing applications for geologic mapping," *IBM Journal of Research and Development* **29,** No. 2, 177–187 (1985).

F. Chatelin, "A probabilistic round off error propagation model," *Reliable Numerical Computing, Oxford University Press,* New York (1988).

P. Dague, P. Deves, and O. Raiman, "Trouble shooting: when modelling is the trouble," *Proceedings of the National Conference on Artificial Intelligence* **AAAI-87,** American Association for Artificial Intelligence (1987).

D. Ho, "A soil thermal model for remote sensing," *IEEE Transactions in Geoscience and Remote Sensing* **25,** No. 2 (1987).

F. Marcotorchino and P. Michaud, *Optimisation en Analyse Ordinale des Donnees,* Masson, Paris (1979).

R. Moreau, *The Computer Comes of Age,* MIT Press, Cambridge (1984), English translation of *Ainsi Naquit L'Informatique,* Dunod, Paris (1981).

O. Raiman, "Order of magnitude," *Proceedings of the National Conference on Artificial Intelligence* **AAAI-86,** American Association for Artificial Intelligence (1986).

Pisa, Italy

C. Bianchi, G. Calzolari, and P. Corsi, "A program for stochastic simulation of econometric models," *Econometrica,* No. 46, 235–236 (1978).

A. Fahmy, L. Panattoni, and E. Todini, "A mathematical model of the River Nile," *Engineering Applications of Computational Hydraulics,* 111–130 (1982).

P. Franchi, C. Paoli, J. Gonzales, P. Mantey, A. Parolo, and J. Simmons, "Design issues and architecture of HACIENDA, an experimental image processing system," *IBM Journal of Research and Development* **27,** No. 2, 116–126 (1983).

C. Paoli, R. Porinelli, A. L'Abbate, and M. G. Trivella, "A computerized autoradiographic technique for simultaneous millimetric mapping of myocardial blood flow and metabolism," *IBM Journal of Research and Development* **30,** No. 6, 627–634 (1986).

C. Paoli, R. Porinelli, C. Lettera, and L. Masera, "Use of a dictionary in conjunction with handwritten texts recognizer," *IEEE Proceedings of the 8th International Conference on Pattern Recognition,* International Association for Pattern Recognition (IAPR), Paris, 699–701 (1986).

M. T. Pareschi, P. Armienti, and G. Macedonio, "A numerical model for the simulation of tephra transport and deposition: Applications to May 18, 1980 Mt. St. Helens eruption," *Journal of Geophysical Research* **93,** No. B6, 6463–6476 (1987).

M. T. Pareschi and R. Bernstein, "Modeling and image processing for visualization of volcanic mapping," *IBM Journal of Research and Development* **33,** No. 4, 406–416 (1989).

Rio de Janeiro, Brazil

M. A. Casanova, F. A. C. Giorno, and A. L. Furtado, *Logic Programming and the Prolog Language* (1987).

R. C. B. Martins and A. V. Moura, *Systematic Derivation of Correct Programs: The Denotational Approach* (1988).

Rome, Italy

F. Antonacci, P. Dell'Orco, and A. Turtur, "A territorial database management system," *Proceedings of the International Conference on Management of Data*, ACM SIGMOD, Orlando, FL (June 1982).

F. Antonacci, M. T. Pazienza, M. Russo, and P. Velardi, "A system for the analysis and generation of sentences in Italian" *Conceptual Graphs for Knowledge Systems*, J. Sowa and N. Foo, Editors, Addison-Wesley Publishing Co., Reading, MA (1989).

F. Antonacci, M. T. Pazienza, M. Russo, and P. Velardi, "Representation and control strategies for large knowledge domain: An application to NLP," to be published in the *Journal of Applied Artificial Intelligence* (1989).

R. Benzi, H. R. Hansen, and A. Sutera, "On stochastic perturbations of simple blocking models," *Quarterly Journal of the Royal Meteorological Society* **110** (1984).

P. Carnevali, A. Selloni, A. Baratoff, and E. Stoll, "Current distribution in the scanning vacuum tunnel microscope: Free electron model," *Journal of Physics C. Solid State Physics* **17** (1984).

S. Cavaliere, M. Fantini, and A. Turtur "An integrated system for printing and publishing applications," *Computer-Generated Images*, Springer-Verlag, New York (1986), pp. 436–447.

E. de Castro, G. Cristini, A. Martelli, C. Morandi, and M. Vascotto, "Compensation of random eye motion in an electronic ophtalmoscope: Preliminary results," *IEEE Transaction on Medical Imaging* (March 1987).

S. Di Zenzo, R. Bernstein, S. De Gloria, and H. G. Kolsky, "Gaussian maximum likelihood and contextual classification algorithms," *IEEE Transactions on Geoscience and Remote Sensing* **GE-25,** No. 6 (1987).

P. D'Orta, M. Ferretti, A. Martelli, S. Melecrinis, S. Scarci, and G. Volpi, "A speech recognition system for the Italian language," *IEEE 1987 International Conference on Acoustic, Speech and Signal Processing*, Dallas (1987).

P. D'Orta, M. Ferretti, and S. Scarci, "Fast speaker adaptation for large-dictionary real-time speech recognition," *IEEE Workshop on Speech Recognition*, Arden House, Harriman, NY (1988).

G. Fronza and P. Melli, Editors, *Mathematical Models for Planning and Controlling Air Quality*, Pergamon Press, Elmsford, NY (1982).

A. Fusi and G. Sommi, "Distributed virtual systems," *Proceedings of the 2nd International Conference on Distributed Computing Systems*, Paris (April 1981).

J. Gazdag and P. Sguazzero, "Migration of seismic data by phase shift plus interpolation," *Geophysics* **49** (1984).

L. Lenzini and G. Sommi, "Architecture and implementation of RPCNET," *Proceedings of the Third International Conference on Computer Communication*, P. K. Verma, Editor, Toronto (1976).

E. Runca, P. Melli, and F. Sardei, "Finite difference and variational methods applied to numerical treatment of advection and diffusion of environmental pollutants from a line-source," *Journal of Computational Physics* **59,** No. 1 (1985).

G. Volpi, G. Gambolati, L. Carbognin, P. Gatto, and G. Mozzi, "Groundwater contour mapping in Venice by stochastic interpolation. 2. Results," *Water Resources Research* **15,** No. 3 (1979).

Tokyo, Japan

J-K. Hong and J. Iisaka, "Coastal environmental change analysis by Landsat MSS data" *Journal of Remote Sensing of Environment* **12,** No. 2, 107–116 (1982).

A. Koide, A. Doi, and K. Kajioka, "Polyhedral approximation approach to molecular orbital graphics," *Journal of Molecular Graphics* **4,** No. 3, 149–156 and 160 (1986).

T. Nishino and T. Fujisaki, "A Stochastic parsing of Kanji compound words," *Transactions of the Information Processing Society of Japan* **29,** No. 11, 1034–1042 (November 1988).

M. Sakaki, H. Samukawa, and N. Honjou, "Effective utilization of the IBM 3090 large virtual storage in the numerically intensive computations of *ab initio* molecular orbitals," *IBM Systems Journal* **27,** No. 4, 528–540 (1988).

M. Udo, Y. Akimoto, S. Kaneko, and T. Sanuki, "Japanese APL language system on IBM multistation 5550," *ACM/SIGAPL APL Quote Quad* **16,** No. 4, 326–334 (July 1986).

O. Wakita and S. Kadekawa "Computer for the blind and the deaf," *Journal of the Information Processing Society of Japan* **25,** No. 5, 462–470 (May 1984).

Y-L. Yang, A. Kondo, and J-K. Hong "An assessment of environmental changes by Landsat data," *Ocean Space Utilization '85* **1,** Springer-Verlag, NY (1985), pp. 715–722.

Winchester, United Kingdom

M. L. Cocklin, "Digital chest radiology: Spatial and contrast resolution requirements," *Proceedings of the International Symposium on Computer Assisted Radiology*, Springer-Verlag, Berlin (1987), pp. 215–219.

H. M. Cole, T. R. Harris, E. W. Stockley, and H. V. Wheale, "Interactive three-dimensional display of reconstructed neurons," *Journal of Physiology*, 400 (1988).

G. Kaye, "The design of the database for the survey of English usage," *Corpus Linguistics, Hard and Soft*, M. Kyte, O. Ihalainen, and M. Rissanen, Editors, Rodopi Publisher, Amsterdam (1988).

S. G. C. Lawrence, B. J. Williams, and G. Kaye, "The automatic phonemic transcription of English," *Orthography and Phonology*, P. A. Luelsdorff, Editor, J. Benjamins Publisher, Amsterdam (1987).

R. W. Phippen, G. R. Luckhurst, and D. J. Adams, "Computer simulation studies of anisotropic systems XVII: the Gay-Berne model nematogen," *Molecular Physics* **61,** No. 6, 1575–1580 (1987).

P. Reilly and R. Zambardino, "A computer-based system for the organization and processing of site-location and boundary data," *Computers and Humanities* **27**, 177–191 (1987).

J. R. Woodwark, "Eliminating redundant primitives from set-theoretic solid models by a consideration of constituents," *IEEE Computer Graphics and Applications* **8**, No. 3, 38–47 (May 1988).

Former IBM Scientific Centers

Bari, Italy

F. Antonacci, P. Dell'Orco, and V. N. Spadavecchia, "AQL: an APL-based system for accessing and manipulating data in a relational data base system," *Proceedings of the APL 76 Conference*, G. Truman Hunter, Editor, Ottawa (September 1976).

P. Dell'Orco, V. N. Spadavecchia, and M. King, "Using knowledge of a data base world in interpreting natural language queries," *Information Processing 77 (Proceedings of the IFIP Congress 1977)*, B. Gilchrist, Editor, Toronto, North Holland Publishing Company (August 1977).

Brasilia, Brazil

M. A. Casanova and A. V. Moura, *An Introduction to Distributed Data Bases Management Systems* (1984).

Grenoble, France

J. Bellino and C. Hans, "Virtual machine or virtual operating system?" *ACM-SIGARCH-SIGOPS*, Harvard University, Cambridge, MA (1973); may be obtained from the Association for Computing Machinery.

J. Rodriguez-Rosell, "Experimental data on how program behavior affects the choices of scheduler parameters," *ACM*, Stanford University, Stanford, CA (1971); may be obtained from the Association for Computing Machinery.

Nguyen Thanh Thi, "A graphic conversational aid for error searching and teaching of an operating system," *ACM-SIGARCH-SIGOPS*, Harvard University, Cambridge, MA (1973).

Houston, Texas

M. M. Adibi, P. M. Hirsch, and J. A. Jordan, Jr., "Solution methods for transient and dynamic stability," *Proceedings of the IEEE* **69**, 951–958 (1974).

T. A. Hughes, et al., *Application of Security Indices and Restraints to a Real Power System, Final Report to the Electric Power Research Institute* (1975).

L. B. Lesem, P. M. Hirsch, and J. A. Jordan, Jr., "Computer synthesis of holograms for 3-D display," *Communications of the ACM* **11**, 661–674 (1968).

L. B. Lesem, P. M. Hirsch, and J. A. Jordan, Jr., "The kinoform: A new wavefront reconstruction device," *IBM Journal of Research and Development* **13**, 150–155 (1969).

George Paul and M. Wayne Wilson, *The VECTRAN Language*, PASC Technical Report G320-3334, IBM Corporation (August 1975); may be obtained through IBM branch offices.

New York City, New York

J. Greenstadt, "On the reduction of continuous problems to discrete form," *IBM Journal of Research and Development* **3**, No. 4, 355–363 (1959).

M. Grigoriadis and W. F. Walker, "A treatment of transportation problems by primal partition programming," *Management Science* **14**, No. 9, 565–599 (1968).

D. A. Quarles and K. Spielberg, "A computer model for global study of the general circulation of the atmosphere," *IBM Journal of Research and Development* **11**, No. 3, 312–336 (1967).

Peterlee, United Kingdom

M. G. Notley, *The Peterlee IS/1 System*, Report UKSC-0018, IBM Scientific Centre (March 1972).

S. J. P. Todd, "The Peterlee relational test vehicle—A system overview," *IBM Systems Journal* **15**, No. 4, 285–308 (1976).

Philadelphia, Pennsylvania

P. G. Comba, "A language for three-dimensional geometry," *IBM Systems Journal* **7**, Nos. 3 and 4, 292–306 (1969).

K. Spielberg, "Algorithms for the simple plant location problem with some side conditions," *Operations Research* **17**, No. 1, 85–111 (1969).

Venice, Italy

G. Gambolati, "Use of over-relaxation techniques in the simulation of large groundwater basins by the finite element method," *International Journal of Numerical Methods in Engineering* **9**, No. 1, John Wiley & Sons, Inc., New York (1975).

E. Runca, P. Zannetti, and P. Melli, "A computer-oriented emissions inventory procedure for urban and industrial sources," *Journal of the Air Pollution Control Association* **6** (1978).

Wheaton, Maryland

A. Hassitt, *Solution of the Stochastic Programming Model of Reservoir Regulation*, Report 320-3506, IBM Washington Scientific Center (May 1968).

H. H. Wang, P. Halpern, and P. Wrotenbery, *Primitive Equation Limited Area Fine Mesh Numerical Weather Prediction Model*, Report 320-3512, IBM Washington Scientific Center (October 1968).

Harwood G. Kolsky *University of California, Santa Cruz (UCSC), Computer Engineering Department, 225 Applied Sciences Building, Santa Cruz, California 95064.* Dr. Kolsky is a retired IBM Fellow, Palo Alto Scientific Center. He is a physicist who became a computer scientist when that field was new. His work has been on a variety of topics including programming languages, scientific applications, and digital image processing. He received his B.S. degree in engineering physics from the University of Kansas and his M.A. and Ph.D. degrees in physics from Harvard University. After seven years at the Los Alamos National Laboratory, Dr. Kolsky joined IBM in 1957 in Poughkeepsie as a member of the product planning group for the STRETCH (IBM 7030) computer. In 1959 he became assistant manager of the IBM Federal Systems Division office in Omaha, Nebraska. Following this he spent some time at FSD headquarters, before being named manager of the systems science department of the San Jose Research Laboratory in 1961. In 1962 he headed an advanced technology group in the Advanced Systems Development Division at Los Gatos, California. He joined the Palo Alto Scientific Center when it was formed in 1964 as manager of the atmospheric physics group. Later he headed projects in programming languages and digital image processing. Dr. Kolsky was named an IBM Fellow in 1969. He served on the IBM Corporate Technical Committee at Armonk, New York, from 1974 to 1975. He was also head of the Board of Consultants for the IBM European Scientific Centers. In 1985 he came to the new UCSC computer engineering board of studies as a visiting professor. In 1986 he retired from IBM and began a new career as a full-time professor at UCSC. He is a member of Sigma Xi, the American Physical Society, and a senior member of the IEEE.

Richard A. MacKinnon *IBM Cambridge Scientific Center, 101 Main Street, Cambridge, Massachusetts 02142.* Mr. MacKinnon has been manager of the Cambridge Scientific Center since 1975. He holds a B.A. degree in American history from Yale University (magna cum laude) and an M.B.A. in finance from the Harvard Graduate School of Business Administration. He joined IBM in 1962, working first as a systems engineer, then marketing representative in a banking territory in Boston. In 1966 Mr. MacKinnon installed the first on-line banking system on System/360. From 1967 to 1969 he worked in the banking development department in the Data Processing Division (DPD) headquarters, White Plains, New York. In 1970 he became systems support manager in IBM's largest DPD region, New England and upper New York state. There he was responsible for introduction and early support of System/370, its virtual storage operating systems, and its associated hardware and system software products. In 1972 Mr. MacKinnon was responsible for the installation of the second System/370 shipped by IBM, which was the first system installed using OS/360. In 1975 he joined the Cambridge Scientific Center as its manager. Since 1976 he has taught the Advanced Computing Systems course at the MIT Sloan School of Management where he currently holds the position of senior lecturer. He is the author of three previous papers in the *IBM Systems Journal* on IBM multiprocessing and virtual machine operating systems. Outside IBM, Mr. MacKinnon is a member of the foundation and the board of managers of the Massachusetts Eye & Ear Infirmary, a Harvard University teaching hospital founded in 1824. He is also on the visiting committee of the radiation medicine department of the Massachusetts General Hospital and an elected member of Phi Beta Kappa.

IBM Scientific Centers map, Reprint Order No. G321-0097.

Paper Reprint Order No. G321-5372.

G321-5372-00

ACTIVITY REPORTS
of
SCIENTIFIC CENTERS
for
NOVEMBER 1987

* *CAMBRIDGE*
* *PALO ALTO*
* *LOS ANGELES*

Cambridge Scientific Center

November 1987 Highlights

KEY ACCOMPLISHMENTS

Application Technology: (CSC)

Artificial Intelligence: (CSC)

Communications and Networking: The TSAF/APPN project activities were presented to a VM/SP Release 7 planning meeting in Endicott. The result is a recommendation that APPN for VM be put into the product plan for 12/88. APPN/VM would come from SPD Toronto. A second release of APPN/VM for 9/89 would include Collection Support based on the Cambridge TSAF/APPN project. This month APPN was declared to be part the SAA architecture strategy. (CSC)

Operating Systems: The Distributed VM Systems Lab has delivered the following programming items to Endicott for inclusion in VM/SP 5.1: (1) code which runs in Remote VM systems to monitor the status of CP-owned file space and automatically initiate corrective action if predetermined thresholds are exceeded; (2) routines to automatically free VM spool file allocations based on a predetermined set of criteria; (3) a monitor for service machines which reinitiates service if problems occur; (4) a facility whereby a Central site can remotely power off 9370s in its network; and (5) programmable operator action routines. (CSC)

Programming Technology: Andy Pierce has made available a suite of demonstration code which shows off X-Windows support for the Megapel adapter and display hardware. These demos have been installed on an IBM conferencing disk to enable marketing reps to demonstrate hardware function to potential customers.

The Real-time Explanation and SuggestiON (REASON) system was presented to 150 attendees of the Expert Systems ITL in San Jose, Ca. Both Donna Lamberti's paper and demonstration received strong support.

The first draft of the REASON architecture document has been delivered to ESD for evaluation, and the project is now moving toward implementation of a prototype suitable for daily use.

The JAWS Windowing System has been successfully ported to OS/2 from PC-DOS. This version of JAWS is currently slated for inclusion in a CPD network configuration product. (CSC)

STATUS UPDATE

Application Technology: (CSC)

Artificial Intelligence: (CSC)

Communications and Networking: Coding and testing of the Bus Controller Microcode for Version 2 of the Cambridge Control Unit were finished. The Bus Controller routes all data traffic between channels and PCs connected to the CCU.
Data transfer between two PC's attached to the V2 CCU was demonstrated using the 3088 (channel-to-channel adapter) interface and achieving a 3.0 MB/sec data transfer rate. (CSC)

IPD Architecture reviewed the status of print server activities from many divisions in IBM at a meeting in Boulder. The review included print servers under development for OS/2 Presentation Manager, the OS/2 Extended Edition LAN Manager, Advanced Function Printing, Enhanced Connectivity Facility and work from Almaden Research and Cambridge. Currently, many of these print servers cannot exchange print requests. The objective of the IPD Architecture organization is to define a means of exchanging print files and print requests. (CSC)

Cambridge Scientific Center

November 1987 Highlights

VNET/BITNET Gateway negotiations have been completed with IIN to transfer the Gateway to that organization. BITNET is a VNET-like network of 300 US Universities. The network is used solely for the exchange of research and educational information. The VNET/BITNET Gateway allows authorized IBMers to communicate and exchange information with academic research partners. (CSC)

Operating Systems: Cambridge has continued work on porting UNIX applications and utilities from 4.3bsd to the Quicksilver distributed operating system under development at Almaden Research.

Cambridge has completed installation of the latest version of IX/370 on a 9370 processor with an ASCII subsystem. This system replaces the IX/370 system on the 4361. (CSC)

Programming Technology: The port of the Megapel based GML foils package to X-Windows is proceeding. This work will also be released on the RTTOOLS disk when complete. The foils package takes advantage of the Megapel anti-alias text support designed and built by Cambridge.

Andy Pierce has worked closely with the Cornell supercomputer project to show the benefits of using the RT to operate cooperatively with the mainframe.

A videotaped presentation of the CCE project was completed by Cambridge staff.

A compiled version of the REASON expert system has achieved a ten-fold improvement in performance in recent tests. (CSC)

NEW OPPORTUNITIES/INVESTMENTS

RACF/VM was installed and made operational on the 3081.

An Information Asset Security class was held or Cambridge location users. (CSC)

Discussions have been initiated with the University of Maine concerning potential joint work as a part of the Distributed VM Systems Lab. Maine would evaluate and investigate the appropriateness of VM/IS software in physically dispersed heterogeneous systems and communications environments. (CSC)

BUSINESS CONCERNS

IBM Cary is moving closer to a withdrawal of the Virtual Machine Performance Planning Facility (VMPPF) program offering. It is uninterested in releasing completed new features of VMPPF which predict the effect of moving from VM-to-VM/XA configurations and recently-announced hardware. Cary is unable to deal with business cases and poor marketing visibility. We are escalating a review of ASD and Cary's handling of this type of software. (CSC)

GENERAL COMMENTS

(CSC)

R. A. MacKinnon, Manager
Cambridge Scientific Center
November 30, 1987

PALO ALTO ACTIVITY REPORT - NOVEMBER

IBM Confidential

Key Accomplishments

Artificial Intelligence

A prototype expert system for Maintenance Marketing has been developed and demonstrated to Messrs. Boler and Timpson of ISG and McDowell of NSD. The prototype was well-received; additional resources have been offered to turn the prototype into a system to be used by field marketing groups. Plans are being put into place to do so.

PASC organized and hosted the 1987 EXpert System ITL conference held at the Almaden Research Center, November 16-20. There were more than 300 IBers from 19 different countries attending this conference. PASC contributed eight papers and four demos to the conference.

Operating Systems The new prototype of MVS/NSP (No Systems Programmer) was demonstrated in Boeblingen, kGermany to the MVS/Entry development group. A working relationship has been established between our groups and we are working towards transferring NSP code to them at the end of 1988 for inclusion in a future MVS product.

Programming

The effective I/O transfer rate is a limiting factor on the performance of many large scale computing problems. Currently, under CMS, we are effective- ly limited to a tranfer rate of about 1.5 MB/sec., although, of course, channel speeds are much higher. Our prototype system for "data striping," i.e., the distribution of data sets across multiple channels, was recently run at the Washington Systems Center utilizing 15 channels. An effective tranfer rate of just under 21 MB/sec. was obtained. We are working with GPD towards inclusion in a product.

Status Update

Artificial Intelligence

KNACK, the knowledge acquisition project at PASC, has reached a major mile- stone: a knowledge base acquired with the KNACK prototype can be converted for ESE consultation on the 370. Specifically, the KNACK prototype can be used to acquire classification-type knowledge using various interactive and batch commands to form concepts and relations, and perform "is-a: consistency on the knowledge base.

:h4Operating Systems As part of the continuing effort to support the Intel study, the PASC team shipped IXX 1.4 as scheduled. In additon to fixes, this release features the support of 200 open files as compared to the support of 20 open files before. This release has been running at Intel for a two weeks without outage. Our next shipment will be designated AIX 0.0 and is scheduled for January. This should be based on product code involved in AIX/370.

Programming Technology

An interim plan to provide initial VS/FORTRAN capability under AIX/370 has been formulated. We expect it to be a joint PASC-GPD plan. Much of the work to modify the H-Assembler to provide a cross development environment to support assembly under CMS has been completed.

Application Technology

Significant progress has been made on completing the Workstation Visualiza- tion Environment prototype. The source code of the communication package has been obtained, and it is being integrated into the Microsoft Windows environment on the PS/2.

Excellent progress continues to be made on the development of an Image Science and Application Workstation. The first version will be on a 5080 under VM/CMS, while later versions are targeted towards independent workstations.

Other

Vector Support Program - A number of PASC people assisted with the conver- sion and set up of four of the codes in the Indiana University benchmark set. Consultation was provided on two codes of the nine code set. This

LOS ANGELES ACTIVITY REPORT - NOVEMBER

IBM Confidential

KEY ACCOMPLISHMENTS

* *Applications Technology*

The Harmony Intelligent Tutoring System (HITS) was presented at the Association for the Development of Computer-based Instructional Systems (ADCIS).

The S*P*A*R*K and STRATEGIST expert systems for the identification of strategic opportunities for Information Technology were transferred to Sterling Forest for use as Expert Systems Environment demos.

"Convexity Algorithms in Parallel Coordinates", by A. Inselberg, M. Reif, and T. Chomut, has been published in the the prestigious Journal of the ACM; the paper is a companion to earlier papers on the mathematical basis of Multi-Dimensional Graphics (MDG).

* *Artificial Intelligence*

With help from Austin, we have benchmarked the robotics edge point image processing code on the RT APC; the code ran in 4.5 seconds vs. 6.0 seconds on the Sun 3/260.

* *Operating Systems*

None.

* *Programming Technology*

None.

* Other

None.

STATUS UPDATE

* *Application Technology*

We have studied the Apple Hypercard in the context of educational workstations. Many of the Hypercard's functions will have to be provided in IBM's offerings, although we will also need to integrate existing PS/2-based tools in ways which are free of the Hypercard's limitations. For example, it appears that the object-oriented programming language ACTOR would provide a good interactive development environment with excellent "hooks" to the services and applications running on the PS/2.

A draft Scientific Center Report on the Multi-Dimensional Graphics (MDG) based air traffic control algorithm has been completed. The draft demonstrates the clear superiority of the new algorithm.

* *Artificial Intelligence*

None.

* *Operating Systems*

CFE/QICS Version 2.0 testing has continued with only minor bugs remaining; we have gained approval to get VTAM source code so that we can work on the VTAM performance bottleneck, the solution to which is expected to improve performance to 750 transactions/second on a 3081KX.

The Communications Front End, taken as a separate package, is being considered as a possible 1988 CPD product offering, as recommended by the DSD High Volume Transaction Processing task force; it is also consistent with the GPD strategy for transaction management.

* *Programming Technology*

None.

NEW OPPORTUNITIES/INVESTMENTS

Research Division is evaluating Multi-Dimensional Graphics (MDG) for the development of a new statistical package, to be done jointly with a customer.

Use of the CFE/QICS technology is gaining support with FSD Gaithersburg for a bid with Bell Atlantic; key factors are a high transaction rate, which CFE/QICS can meet, and low cost/transaction, which has yet to be determined.

BUSINESS CONCERNS

IBM's expert system strategy views the PS/2 as primarily an intelligent front end to a host. The features of ESE 2 would go far toward meeting the needs of Intelligent Computer Aided Instruction (ICAI) software developers if available in a stand-alone PS/2 version.

FSD is bidding on a large contract with DOD, a part of which is an English/Korean machine translation system. FSD expects to vendor the MT system rather than try to assemble a team from within IBM. Assuming the predictions on the value of MT (a multi-billion market into the 21st Century) are correct, a concentrated effort by knowledgeable IBMers might propel IBM into a profitable leadership position in MT.

There is little visible marketing of the IBM Music Feature, one of the few state-of-the-art IBM offerings in the "hypermedia" arena. This leads to more general concerns on IBM's ability to market to the non-traditional end-users to whom hypermedia is a major attraction.

Scientific Centers in IBM: a History

Harwood G. Kolsky

Richard A. MacKinnon

TABLE of CONTENTS

Dedication: Dr. Louis Robinson, Ph.D

1926 – 1985

Dr. Louis Robinson, Ph.D.

1926-1985

What follows in these pages is dedicated to Dr. Louis Robinson. He enjoyed a most diverse and colorful career in IBM which upon two separate occasions found him directly responsible for Scientific Computing and the Scientific Centers in the United States. As a result his contact with the Scientific Centers was the longest of any IBM executive.

Lou Robinson brought his most effective and inspirational talents to the Scientific Centers. Early in his career he worked with Drs. Charles DeCarlo and Herman Goldstine in the formation and embryonic phases of the Scientific Centers. In 1955 he founded and managed the Mathematics and Applications Department in New York City. This group became the New York Scientific Center in 1964. In 1975 he returned for the second time as Director of Scientific Computing and had direct responsibility for the US Centers.

It was in this final tour he helped strengthen the centers and their contributions to scientific and engineering computing within IBM while fostering a renaissance in out-reach through university joint studies. His liaison with the IBM Chief Scientist and the Corporate Office was effective, protective, courageous, and visible. It enabled the Scientific Centers to survive over the years and contribute in their historic role as an IBM corporate resource.

Lou was a graduate of the University of Massachusetts (M.S., summa cum laude) and received M.A. and Ph.D. degrees in mathematics from Syracuse University, where he taught before joining IBM in the Data Processing Division in 1953.

His tenure in IBM saw him intimately involved with scientific marketing, measurement and evaluation of computer performance, and early advanced computer applications such as Braille translation and its printing. He managed IBM's effort for the NASA tracking of the Vanguard satellite. He was Director of Standards in the Data Processing group from which position he served on the body which adopted the Universal Product Code (UPC). In his last assignment (1982-85) he was Corporate Director of University Relations.

Lou Robinson became well-known for his diverse efforts to explain computers and computing. This ranged from his Encyclopedia Americana chapters "*Computers*" and "*Data Processing*" to his consulting with TIME Magazine on its cover story "*The Computer Society.*" He appeared frequently on national television programs and internally in IBM video tapes. He advised the Charles Eames Office on their exhibition at the Metropolitan Museum of Art, and consulted with EPCOT Center on computer capabilities.

Professionally, Lou was a national Phi Kappa Phi Fellow. He served as a national lecturer of the Association for Computing Machinery, and held memberships in the Institute of Management Sciences, American Mathematical Society, New York Academy of Science, Mathematical Association of America, and the honor society, Sigma Xi. He also held honorary doctorates from Michigan Technological University and Maryville College.

The Scientific Center community took great pride in Lou's accomplishments and, in view of his long bout with cancer, marvelled at his energy and on-going good cheer. He never seemed to have a dark moment. He enjoyed both a uniquely long perspective on the past and gifted vision of the future of computing. He combined them with buoyancy and cheer, and when willing to reminisce, displayed impish irreverence both toward himself and those who took themselves too seriously. His ability to drop everything and help anyone -- employee, colleague, or just a fellow human -- made him a legend to those touched by his grace and thoughtfulness.

At Lou's funeral the rabbi characterized him as a man who left many projects underway and undone at his untimely death in 1985. At the same time he was memorialized as a man with so many enthusiasms and responsibilities that this would have true at any point in his life.

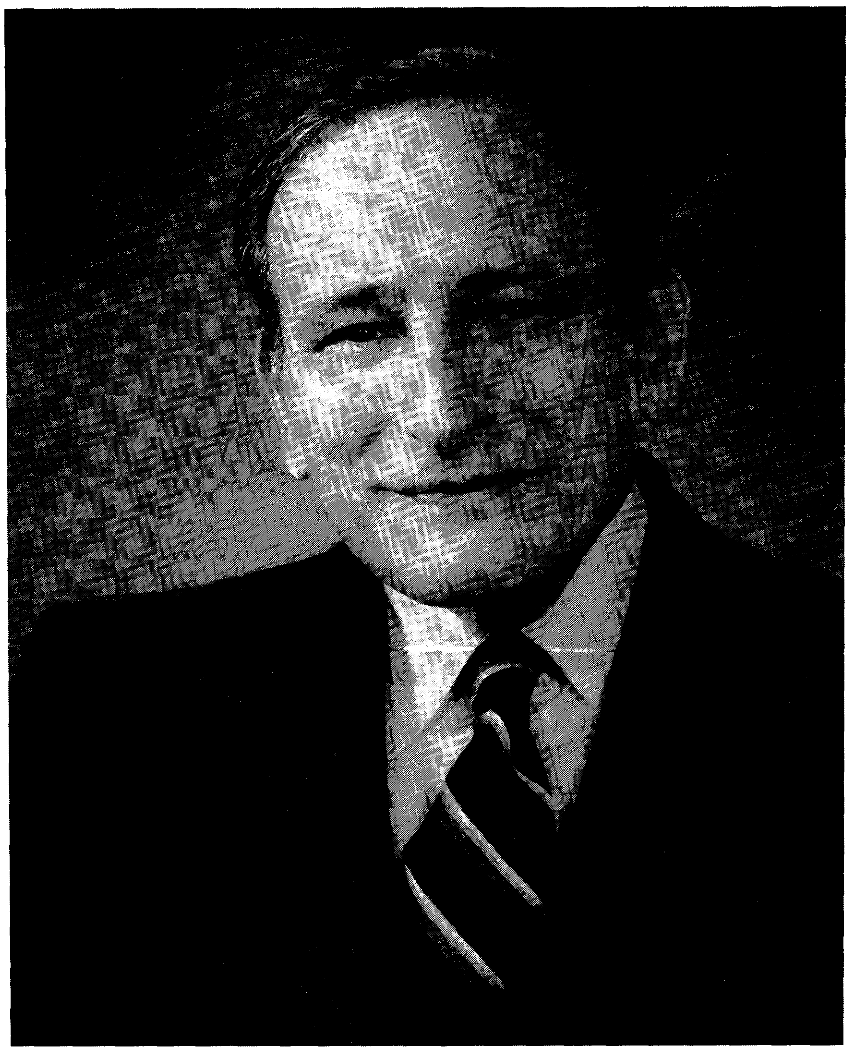

ABSTRACT & FOREWORD:

About This Paper

Scientific Centers in IBM: a History

Harwood G. Kolsky

Richard A. MacKinnon

Foreword - About This Paper

In 1989 the IBM Scientific Centers celebrated their twenty- fifth anniversary. They were formed in 1964 to fill a real need within IBM to match the rapid changes in technology at the time with correspondingly rapid responses. The concept has been so successful over the years that there are now 17 Centers world wide operating as autonomous organizations, providing IBM with the ability to respond rapidly to the evolution of computer technology for IBM and for its scientific customers.

From the beginning the primary mode of operation has been establishing long range contacts with leading scientific customers, understanding their problems, defining solutions, developing prototypes, and insuring IBM's responsiveness to their needs.

During this quarter century the Centers have provided technical leadership in almost every branch of computer science. Many individual accomplishments within the Centers have made a large impact on impact on IBM, not only in terms of the product line, but also in terms of upgraded scientific and technical quality for other projects within IBM, basic research, market support, and a generally enhanced relationship with scientific customers. Thus, technology transfer remains the Centers' key contribution for IBM.

The report contains: an introduction to the mission, scope, and history of the Centers; a brief description of each Center's history, and accomplishments; a selected bibliography for each Center; and a historical appendix on management issues. Portions of this report were used as the basis for the 1989 issue of the *"IBM Systems Journal" (Vol. 28, No.4, 1989)* article entitled: *"History and Contributions of the IBM Scientific Centers"* by the authors. [Ca10] [1]

Unlike the earlier article, this report provides an ***Historical Appendix*** to give greater insight into the management challenges of the United States Centers as well as updating the material to reflect the unfortunate closing of Los Angeles (in January 1990) and Cambridge/Palo Alto (in July 1992). With the exception of discussion of the closing in the ***Afterword,*** descriptions of each Scientific Center and its work remain as originally written in 1989 for the 25th anniversary.

---

[1] We gratefully acknowledge the permission of "*IBM Systems Journal*" Editor, Mr. Gene Hoffnagle, for permission to use portions of the referenced IBMSYSJ article.

Scientific Centers in IBM: a History

Scientific Centers in IBM: a History

The IBM Scientific Centers were chartered, and the first four established, in 1964 to provide IBM with the ability to respond rapidly to the scientific marketplace and to changes in technology.

The idea of having a small, entrepreneurial organization within a much larger company, has appeared again and again in American industry. "Skunk works," "ad tech groups," and "back room projects" are a part of American corporate folklore. The IBM Scientific Centers were created to fill just such a role for IBM. The emphasis was and is on having an autonomous organization of highly skilled professionals who can develop scientific solutions and applications at the leading edge of technological change.

Consider the impact of these Center contributions:

- VM and CMS (now known as VM/SP), VNET and BITNET, the Vector and Parallel FORTRAN compilers for the IBM 3090, the IBM 7350 Image Processing System (formerly Hacienda), the APL Language and Products based on it, the first IBM PC Prototype (IBM 5100), the ESE Expert Systems, the AIX/370 Distributed Operating System, and the Deaf Children project.

- The Centers have worked in almost every branch of computer science -- computer architecture, numerical analysis, algorithms and data structures, operating systems, relational data bases, languages, compilers, microcode, networking, artificial intelligence and knowledge-based systems.

- The Centers have worked on an incredible list of computer applications from basic physics, chemistry and mathematics to agriculture, oceanography, surface hydrology, to management science, econometrics, operations research, and commercial packages. They have had joint studies with dozens of universities.

- The Centers have been heavily involved in human issues, such as the medical applications of computers, speech analysis and synthesis, computer vision, and in natural language projects in Spanish, Catalan, Arabic, Italian, French, German, Chinese and Japanese.

- The Centers have on a world-wide basis been involved in problems of national importance to dozens of countries, such as regional planning, food and oil production, pollution studies, mapping the Nile, and "saving Venice."

Origins:

When the first IBM Scientific Centers were founded twenty-five years ago in the United States, they were built on already firmly established traditions. The IBM Research Division was well-established, as were several large product development laboratories. The Applied Science Department in New York City under Dr. Cuthbert Hurd had been very successful in working with scientific customers during IBM's early expansion into electronic computing in the mid-1950s.

Key models for the Centers were the joint IBM and University of California at Los Angeles (UCLA) Western Data Processing Center (WDPC) located on the UCLA campus, and a similar Center near the Massachusetts Institute of Technology (MIT), both of which had been established in 1957. Much of the philosophy and mode of operation for the new Centers came directly from these two prototypes, as did many of the original staff.

The founder of both the earlier centers and the IBM Scientific Centers was Dr. C. R. "Charlie" DeCarlo, who was on the staff of IBM senior executive Arthur K. Watson. Dr. DeCarlo, a charismatic leader, had headed the original Applied Science organization.

In 1960 the specialized Applied Science group was replaced by the more general Systems Engineering organization. However, the need for the Applied Science function still existed in the field. Many people probably viewed the new Centers being proposed in 1964 as just another DeCarlo reincarnation of Applied Science. However Arthur K. Watson foresaw a continuing need for an IBM commitment to applied science and gave his full support to the endeavor. The result, after 25 years, is the present international network of Centers.

The Planning Conference:

A planning conference for the Systems Research and Development Centers (as the Scientific Centers were first called) was held in Los Angeles February 10-14, 1964. The report of the conference indicates, by its great detail, that much staff preparation had been done beforehand in 1963. There were three sections: (1) the current state of the scientific market, (2) a proposed general plan for supporting the scientific market through the centers, and (3) suggested operating policies for each of the centers.

The stated General Strategy for the Centers was *"to provide within IBM the ability to match rapid change with correspondingly rapid response. The primary mode in which this function will be performed is by establishing long range contacts with leading scientific customers, understanding their problems, defining solutions, and insuring that capable IBM agencies respond to the need."*

In the Operations section the report said: *"The major portion of SR&D activity will consist of operations in and about approximately seven SR&D Centers located in areas of concentrated scientific endeavor. Each center will be staffed with a number of professionals in disciplines commensurate with the objectives of the center."*

"SR&D should work toward long standing professional relations which can continue through fluctuations in the local sales atmosphere. It is clear, then, that established professional standing should not be compromised for short term reasons."

The key statements in the document emphasized the systems orientation of the Centers and their role in exploring non-strategic prototypes: *"SR&D's contribution in the area of application development will be the development of prototype systems as opposed to officially supported application programs. Once the level of refinement is determined by SR&D through mutual evaluation with the scientific customer, such documentation and systems refinements as necessary will be provided through normal channels."*

This last paragraph became the basis of many important university interactions, future prototypes and joint studies with scientific customers.

The Mission in Practice:

An original objective of the Centers was to have each one involved in the professional and academic environment of a nearby university -- working with professors, involving students in projects, and contributing to the scientific milieu of the university. In practice each Center has also operated on a national and international basis, participating in projects with governments and remote partners. It was also an objective that the Centers influence IBM's technology direction and product offerings. Both objectives have been consistently achieved over the history of the Centers.

Their university setting permits Centers to hire students as employees to work on projects. Students work on the normal technical work, and there are numerous instances of students making significant contributions to a project.

In a typical U.S. Center 70 percent of the staff have advanced degrees. Many of the professional staff have Ph.D.'s, although not necessarily in computer science. (In fact the first Ph.D. in computer science appeared about the time the Centers were founded, so we and the field have grown together. See also *"ODE to a BLACK BOX"* at the end of this section).

During this first quarter century, there have been over 600 individuals at twenty-six IBM Scientific Centers worldwide who have been members of the scientific and support staffs. At the same time a greater number of students, temporary assignees, visiting scientists, and joint study partners have also participated. It is clearly impossible for us to properly recognize the work of all these dedicated individuals in this report. The following section gives a brief statement as of 1989 of the major accomplishments and areas of research and development addressed by the dedicated and diverse groups at each Center. References are made to a bibliography of selected work from each of the Centers.

The authors acknowledge their considerable debt to all who contributed to these descriptions and particularly to the Center managers who so willingly contributed materials at our request.

ODE TO A BLACK BOX [2]

When you and I were boys, men, and all the world was young,
The power of the keypunch was on everybody's tongue.
The 405's would clank and bang. The 513's would purr.
And the gallant little punch card would do all we asked of her.

And Ah! but life was simple. And Oh! but life was gay,
When the glamor of the crossfoot was enough to save the day.
So curses on the twisted brain which first conceived the hunch
That what the world was lacking was a calculating punch.

> *Make it bigger. Make it faster. Make it smarter. Paint it black.*
> *Give me yet another gadget. Get this monkey off my back.*

The relay, it was mighty fast. The tube was faster still.
And a deftly wired plugboard almost always filled the bill.
But wiring was tedious and thinking was a chore,
And the one thing that could help us was a program we could store.

> *Make it bigger. Make it faster. Make it smarter. Paint it gray.*
> *With another little gadget, I'll be sure to save the day.*

But before one stores a program, one has to code it, lad.
And coding was a burden that would drive a sane man mad.
It was clear the one thing needed to make us all ecstatic
Was a simple little program to make coding automatic.

> *Make it smaller. Make it faster. Make it smarter. Paint it red.*
> *With the right assembly language, I'll be sure to knock them dead.*

But assembling's not the final word. Compiling's not much better.
And an "operating system" is a canker and a fetter.
The things we need to save us, (and here's a thought sublime),
Are the concepts and the kludges which will let us share the time.

> *Make it smaller. Make it slower. Make it smarter. Paint it blue.*
> *And this one last little gadget will be all I'll ask of you.*

Now you and I are men, boys, and all the world is old.
And the gadgets that surround us are both talented and bold.
It is time to think of systems. It is time to plan ahead
Else the world will say about us that we should'a stood in bed.

> *Make it smaller. Make it slower. Make it dumber. Paint it black.*
> *'Cause this weird array of gadgets are like monkeys on my back.*

2

by Michael S. Montalbano, IBM SR&D Poet Laureate. Written in honor of the birth of the Palo Alto Center, 1964. [Mike Montalbano was an original member of the Palo Alto Center and served many years as its expert in management science and in the applications of APL system. Although he worked in several other groups after leaving Palo Alto, he was always invited back whenever there was an important event to be immortalized by one of his masterful poems. The last occasion was August 1986 at the retirement from IBM of his old manager, Harwood Kolsky.]

ODE TO A BLACK BOX (Concluded)

Let us develop SYSTEMS, bend the gadgets to our will.
If we do RESEARCH in systems, then a gadget's like a pill.
If you take it when you shouldn't, it will make you mighty sick.
It it's what the doctor ordered, it is sure to heal you quick

Make it bigger. Make it faster. Make it smarter. Paint it ochre.
For a gadget IN A SYSTEM is like four aces and a joker.

The United States

Scientific Centers

THE UNITED STATES SCIENTIFIC CENTERS

Establishment

Four Centers were established in the United States in January 1964, with two more planned. The Los Angeles (California), New York City (New York), and Cambridge (Massachusetts) Centers under the leadership of Dr. Homer H. Givin, Dr. Theodore I. Peterson and Mr. Norman L. Rasmussen, respectively, were drawn from existing or associated groups. An entirely new center, headed by Mr. John W. Luke, was established at the Los Gatos Laboratory in California and moved to Palo Alto, adjacent to Stanford University, in 1965. The other new centers mentioned were be in Chicago and Washington, DC. There was also a headquarters manager and staff at White Plains, headed at first by Dr. Charles R. DeCarlo.

The Centers were transferred from Corporate headquarters to the Data Processing Division on March 25, 1965. Dr. DeCarlo argued that they should be in a line organization to take the leadership role he envisioned for them. Mr. Frank Cary, then president of the Data Processing Division, agreed and took the centers along with Dr. Herman H. Goldstine from the Research Division to lead the new organization. The Corporate Mission of the centers in dealing with scientific customers was emphasized in the transition.

Evolution of the US Centers:

Detailed snapshots of the Centers' activities in the early years are preserved in their many planning documents. Most of the project areas discussed are still key opportunities for IBM today.

The Centers prospered during their first few years. The number and size of projects increased rapidly, as did the number of staff members. There were plans to open several more centers. However, the "1970 Recession," which really started in 1968 and lasted until 1973, had a drastic effect. For example, the Chicago center was "stillborn," although it appeared on many charts and planning documents between 1966 and 1969, complete with project areas and a building location.

Another event which had an important effect on the Centers was the "Unbundling" decision in 1969. Unbundling marked a major change in which IBM began to market software and services separately from hardware. Software products then slowly began to take on their true worth in the industry.

Centers in Washington and Houston were opened but then later closed as "economy measures." By 1973 the Centers had been cut by one-third from their peak in personnel. Of course the Centers were not the only organizations in IBM who had to justify their existence during those hard times.

After Dr. Goldstine the US Scientific Centers had many leaders, including Dr. John C. Porter, Dr. Bruce Gilchrist, Dr. Joseph F. Mount, Norman L. Rasmussen, and Peter Robohm. But it is Dr. Louis Robinson who must be given much of the credit for guiding the Centers during the critical years from 1975 to 1982. Throughout this period, although battling cancer, he articulately and steadfastly promoted the importance of science and the Centers for IBM and its customers.

Richard A. MacKinnon was acting director for several months in 1982. Richard V. Bergstresser took over as leader in December 1982 and maintained the integrity of the US Centers through a series of organizational changes until Los Angeles closed in January 1990. (See ***Afterword : a Final Summing Up*** for further discussion of the fate of the United States Centers 1990-1992.) Dick Bergstresser took another IBM assignment in September 1990, and was the last official Director of U.S. Scientific Centers.

Cambridge, Massachusetts (Established 1964)

The IBM Cambridge Scientific Center (CSC) was founded in 1964 as one of the first Scientific Centers. It evolved from a joint project with the Massachusetts Institute of Technology (MIT), which had produced the Compatible Timesharing System (CTSS) for the IBM 7094. Mr. Norman L. Rasmussen was the first manager. Dr. William P. Timlake was the second, and Mr. Richard A. MacKinnon is the current manager of Cambridge.

While the joint project office was established on MIT premises, the Scientific Center was inaugurated at 545 Technology Square in the same building which also became the home of MIT's Project MAC, later the Laboratory for Computer Science. From its start and continuing through today, CSC has focused machine architecture, interactive computing, programming languages, and systems performance. The early CTSS work clearly influenced the Center's interactive and virtual machine work and clearly gave the Scientific Centers both their first major and most illustrious contribution to IBM and to computer science.

The Virtual Machine/370 (VM/370 System): Much has been written about this unique and widely used system [ca01". Virtual machines first became available in 1968 through IBM's Type III program library under the name "Control Program Model 67 Cambridge Monitor System" (CP-67/ CMS). Norm Rasmussen led this effort and four of his staff -- Robert J. Creasy, Robert J. Adair, John B. Harmon, and Les W. Comeau -- have received principal recognition from IBM as innovators of the virtual machine concept.

CP-67/CMS became an official IBM control program after the software "unbundling" announcement in May 1969. Virtual storage was subsequently introduced in 1972 across the entire System/370. VM/370 became one of the four demand-paged operating systems utilizing the common hardware architecture of System/370. With the interactive computing facilities of Conversational Monitor System (CMS renamed) and a wide variety of applications, the VM/370 system has become the most widely used operating system on the S/370 architecture [ca02, ca03".

VNET: Almost as prominent and significant to CSC was its launching of VNET [ca04". "VNET" is the informal name for the main internal network of the IBM Corporation. Its formal name as it is embodied in the VM system is Remote Spooling Communications Subsystem (RSCS). VNET is a peer-to-peer, store-and-forward network protocol, built using virtual machine principles. It was invented at CSC by Edson C. Hendricks of CSC and Tim C. Hartmann of Poughkeepsie.

With collaborative efforts over many years (starting in 1972), Cambridge worked to implement RSCS and make it available as a VM product (from 1975-79). It is now the major network facility used by IBM for internal communications, connecting over 3400 network computers world-wide.

BITNET: CSC was instrumental in assisting the startup of a major US university network known as BITNET. The network was founded in 1981 as a collaboration between Yale University and City University of New York, and uses RSCS protocols on IBM and other systems [ca05".

BITNET has grown to include approximately 441 institutions on 1622 nodes. Similar networks have arisen outside the United States using these same protocols. European Academic and Research Network (EARN), NetNorth (Canada), and analogous academic networks in the Far East and Latin America comprise an estimated 686 institutions and 1,036 nodes. This networking phenomenon is probably Cambridge's most pervasive and visible contribution.

Multiprocessor VM/370: In 1976 CSC undertook an advanced development role with Poughkeepsie to implement VM/370 in a multiprocessor configuration for IBM Model 158/168 Attached Processors [ca06". This critical work was accomplished in the record time of ten months and was transferred directly for immediate announcement, release, and maintenance. It supported VM/370 on IBM's largest S/370 and set the scene for the Symmetric Multiprocessing (MP) of the 303X and 3090. This is an excellent example of experimental work becoming a product and of successful transfer and partnership with a product division.

VM Performance - Modeling and Scheduling. Two state of the art contributions from Cambridge have benefited the VM system and VM customers over many years. Dr. Yon Bard has long-been involved in various aspects of measuring and modeling the performance of VM systems. The VM Performance Planning Facility (VMPPF) Program Offering was released in 1983 and provides systems engineers and customers an accurate tool for assessing and predicting the effect of configuration changes on a VM system workload.

In 1976 Lynn Wheeler's "Fair Share" Scheduler for VM appeared as the "VM Resource Management Programming RPQ" and had dramatic impact on the responsiveness and throughput of many VM systems. This work became an integral part of all VM systems at a time when VM performance enhancements were much needed, especially for high-end systems. Bard and Wheeler's contributions both travelled the familiar Scientific Center road -- experimental work followed by extensive internal IBM usage for many years before eventual product release.

MVS Preferred Guest Virtual Machine: Another contribution to VM/370 on high-end IBM processors was the availability in 1981 to allow the MVS operating system to run under VM as a "preferred guest." This work has been carried on in the Processor Resource/Systems Manager(TM) (PR/SM) on the IBM 3090. The success of MVS in conjunction with VM/370 owes much to the pioneering work undertaken by Ron Reynolds in 1978-80.

Yale ASCII Terminal System: The Yale ASCII terminal system is an excellent example of significant work undertaken jointly with a university. CSC's collaboration with Yale University Computing Center and IBM Boca Raton provided a way to permit full duplex ASCII capability for the System/370 through a Series/1 front end, without changes to host programming [ca07".

Significantly, a succession of front-end processors (IBM 7171 in 1983 and IBM 4994 in 1984) evolved from the Series/1 for the dedicated purpose of running follow-on versions of this pioneering software.

IBM PC Projects: CSC's ASCII communications work led directly in 1980 to work on the IBM Personal Computer during its initial planning and development. In August 1981 when the IBM PC was announced, a companion communications package the "IBM Asynchronous Communication Support" from Cambridge was also announced. Subsequently the Center also authored the IBM 3101 Emulation Support (1982).

Since 1981 Cambridge has participated in a variety of product efforts at varying levels, ranging from invention and first prototypes of PC/370 in the IBM/XT (1983) [ca08", to the design of a chord keyboard, an enhanced PC keyboard and work on the IBM Convertible PC by IBM Fellow Nat Rochester in 1985.

Small systems support: Starting in 1979 the Center undertook experimental work in the area of distribution, operation, and support of small System/370s. Unattended and remote operation was an early objective and a constant theme, including remote initial program loading, power-off, auto-power-on, and auto restart upon restoration of power. The IBM 4300 series was the earliest beneficiary (1981), the latest are all the 9370 processors (1986) and the AS/400 (1988). CSC's Personal/CMS was introduced with the PC/XT (1983).

Operating Systems: Cambridge has been involved with the UNIX(TM) operating system as long as any laboratory within IBM. This led to collaboration on the VM/IX programming product. In 1988 CSC completed product work on X-Windows for VM, and in 1989 on the MVS operating system. This emphasizes the importance of CSC's long-standing interest in workstation interfaces and its proximity to MIT.

Text processing: CSC's long-standing involvement in text handling is evident since SCRIPT was an early part of CP-67/CMS. After many iterations it is today's Document Composition Facility (DCF). On VNET, it helped pioneer electronic mail in IBM. "IBM Usability Aid" (QPRINT) in 1984 greatly enhanced the usability of the IBM 6670 laser printer. CSC's interest in text continues in efforts for laser print servers on workstations and the Math Formulator product announced in 1988 as a component of the PC DOS 4.0 Windows Kit.

Communications control: CSC's early VNET work focused on wide-area communications. An ambitious 1984 hardware project to attach local area networks to S/370 channels was called the Cambridge Control Unit (CCU). It led to the invention of a new IBM channel architecture called Continuously Executing Channel Interface (CETI) in 1986. The CCU also served as the model for the IBM 8232 PRPQ implemented by Yorktown Research and the Raleigh Laboratory.

POLITE: POLITE was an early effort on WYSIWYG (*"What you see is what you get"*) text system for S/370 and PC. It had a long-term effect at Cambridge. In 1978, POLITE on a early PERQ workstation introduced the Center to local area networking in the form of Ethernet and triggered early investigations of issues of LAN performance and attachment of LANs to System/370.

The user-friendliness of POLITE pioneered the concept of "UNDO" and "REDO" for the backtracking capability of IBM's DisplayWrite/370 text product in 1985. It also influenced CSC's expert and knowledge-based systems prototypes, and it stimulated an on-going interest in advanced hardware displays with windowed and graphical interfaces. POLITE's key impact was less on products than on the technical direction and intellectual life of the Center and its people.

Retrospective Note: In his final report for the Cambridge Scientific Center Richard A. MacKinnon, Center Manager, has updated the Center's accomplishments through 1992 and commented extensively on his perspective of the Center's contributions to IBM. This report is generally available from IBM. See [ca11].

Houston, Texas (1966 - 1974)

The IBM Houston Scientific Center formed in December 1966 was headed by Dr. Joseph F. Mount. Subsequent managers were Dr. M. Stuart Lynn and Dr. James A. Kearns. The Center was closed in 1974 and most of the staff and projects were merged into the IBM Palo Alto Scientific Center.

Much of the work at Houston was focused on the use of array processors, such as the IBM 2938 and 3838. These studies included numerical precision, seismic applications, and development of array algorithms and languages. Other projects were optical and holographic image processing and display, medical data processing for the heart and X-rays, chemical engineering real-time process control, and electrical power systems analysis.

Array processing and signal analysis: From its beginning the Houston center was heavily involved with the hardware, software and applications of the IBM 2938 Array processor and its successor, the IBM 3838. This included numerical precision studies, array algorithms, the development of a seismic applications manual. "VECTRAN", an experimental language for vector and matrix array processing was first developed at Houston [ho01". VECTRAN later had a significant effect on the array handling capabilities of the FORTRAN 8X standard.

Optical and holographic image processing: Projects included the digital construction or encoding of holograms of 3-D objects which only exist in the computer. In the laboratory techniques of optical laser holography and digital computation were brought together [ho02, ho03". A number of applications of filtering and display were studied. The work on holography at Houston was truly pioneering work and IBM secured a number of basic patents in the field.

Medical data processing: The project was primarily concerned with the analysis of electrocardiogram (EKG) signals. Segmental EKG analysis was done on the IBM 1800. Medical X-rays and the modelling of the heart and blood flow in vessels were also studied.

The techniques of optical digital holography were applied in the ultrasonic domain, and images of soft tissue were reconstructed from acoustic wavefronts reflected from the tissue.

Process analysis and control: A number of chemical engineering problems of interest to the oil industry were undertaken, for example, the separation processes in refinery distillation columns.

Power Systems Analysis: In close interaction with IBM Public Utilities Industry, this work led to applications of sparse matrix and variable time step integration for transient stability. The project was moved to Palo Alto in 1974 with emphasis on real-time analysis tools for power systems operations, databases, and graphical distribution facilities management [ho04, ho05".

Los Angeles, California (Established 1964)

The Los Angeles Systems Research Development Center (now LASC), under the leadership of Dr. Homer H. Givin, was one of the original centers founded in January 1964. It was formed from the IBMers assigned to the joint IBM/UCLA Western Data Processing Center (WDPC) on the campus of the University of California at Los Angeles (UCLA).

Starting in Kirkeby Center, LASC has moved several times. From 1984-89 it was on Wilshire Blvd., Brentwood. Recently it moved to Colorado Place in Santa Monica. Past managers of LASC have been Dr. Lewis E. Leeburg and Dr. Bernard D. Rudin. The present manager is Dr. James A. Jordan, Jr.

In the decade prior to 1964, WDPC was a primary education facility for academic computing and for IBM on the West Coast. In the early years, LASC continued WDPC's emphasis on management science, education, engineering, and science. In recent years, some of the projects still focus on those areas, and new areas have been added, as described below.

Direct numerical control: An early project, which started LASC's continuing involvement with Los Angeles' manufacturing customers, involved the direct control of machine tools connected remotely to a computer. Software was developed to cut metal test pieces according to real customer designs. The project was ultimately transferred to an IBM industry group [la01".

Educational information systems: A major early project was undertaken in computer aided instruction (CAI). LASC established a computer center on the new campus of UC Irvine, developed an interactive supervisor for the S/360, and sponsored a number of projects in the development and use of CAI.

This theme continues today in joint work with the UCLA Medical school on a workstation to support clinical instruction and practice. It is also providing general tools for object oriented programming.

Oceanography: A joint research project with the Scripps Institution involved the installation of an IBM computer on board a research vessel. Problems of processing and storing oceanographic data on line at sea were studied [la02".

Geophysics: In a project with Cal Tech, LASC undertook theoretical studies of earthquake structure using computer simulation of complex models of the earth [la03".

Natural language processing (NLP): The first LASC project in this area was a study of medical records. Natural language information storage and retrieval required the development of syntactic and semantic methods [la04".

NLP is still an active area of research at LASC. Applications being investigated include text selection and understanding of structured documents, (such as the "Wall Street Journal") and machine translation systems capable of translating from many source languages to many target languages. There have been important results in lexicons, parsers, knowledge representations, discourse structures, and NLP systems architecture [la05".

Systems and end-user projects: Over the years the project emphasis at LASC has shifted toward systems and end-user projects, while retaining an interest in engineering and scientific topics. For example, operations research techniques were applied in many areas. A key result was a distribution system simulator which was distributed as a product [la06".

LASC's studies of computer-aided design led to the formation of IBM's West Coast CAD/CAM development activities [la07, la08". Earlier studies of the development process were the background for LASC's present application development environment activities [la09". Studies of database security led to improvements in IBM products [la10, la11".

Studies on distributed computing were sparked by Rita Summers' paper on the architecture of networks of resource sharing machines, later translated into the RM system [la12". This paper proved influential in the development of personal computer based local area networks.

Multi-dimensional graphics: A new approach to an old mathematical question of how to visualize geometrical objects in higher dimensions led Dr. Alfred Inselberg to discover some very practical displays for exploratory statistical analysis and algorithms for air traffic and robotics control. Joint studies have been conducted with several universities, in particular UCLA, USC, RPI and George Mason University [la13, la14, la15, la16".

Engineering and scientific systems: Projects have included: exploration of workstations for meteorologists, chemical engineers, biologists, and for general engineering; studies in large scale scientific computing architecture and applications, and in a hardware/software prototype parallel data store [la17, la18, la19, la20".

Robotics Control Systems: LASC's current emphasis is on vision systems for mechanical parts identification and location. Studies are being conducted to use artificial intelligence techniques for computer-integrated manufacturing (CIM) by the Los Angeles-based Institute for Manufacturing and Automation Research (IMAR) [la21, la22".

Enterprise-wide information management: This effort, led by Marilyn M. Parker, is concerned with improving the effectiveness of business systems planning, and integrating it more fully into corporate planning, budgeting, and decision making. A world-wide network of customers and universities now uses some of the results of an LASC joint study with Washington University of St. Louis [la23, la24".

Strategic Information Management Support: A LASC prototype expert system, Business Strategy Advisor (BSA), has been transferred to the IBM Santa Teresa Laboratory to assess competitive software products. In other work, business modeling and a prototype expert system is helping to identify opportunities for information technologies in enterprises [la25".

Visual Programming: Visual Programming uses the familiar task of filling in forms as a means for user/system communication. Stylized form headings are used as visual representations of data structures (e.g. heirarchical data or nested tables). The operations are applicable to relational data as well. Most important, the programs themselves are also represented as forms. The output is "What you sketch is what you get." Data transformation and restructuring is done automatically. Nan Shu's book on "Visual Programming" is a timely contribution to this emerging field [la26".

Security: LASC has studied tools to enhance the security of IBM's internal computer network (VNET) from the execution of potentially destructive files. A prototype expert system for security audits is being tested [la27".

Logic Programming: Many of the current LASC projects take advantage of the power of expression found in the Prolog logic programming language. To improve the portability of these applications, and perhaps their performance, a Prolog to C translator (C-LOG) has been written.

New York City, New York (1964 - 1972)

Founded in 1964, the IBM New York Center was one of the first four Centers chartered. It grew out of the the Mathematics and Application Department which was formed in New York City in 1955 by Dr. Louis Robinson, and later headed in White Plains by Dr. John C. Porter. In 1965 it was renamed the New York Scientific Center with Dr. Theodore I. Peterson as manager. It was later managed by Samuel M. Matsa. There were two basic areas of concentration - management science and engineering design.

The Center was merged with the Research Yorktown APL group and became part of the foundation of the new Philadelphia Scientific Center in 1972.

Management science: New York's work in this area included the development of new computational programming, simulation, and network analysis. New applications were in such areas as aircraft and airline crew scheduling, solving various scheduling problems for railroads, and solving plant location and machine sequence problems [ny01, ny02, ny03".

Engineering design: The engineering design group designed a Special Control Unit to facilitate the attachment of non-standard data acquisition equipment to the System/360. They devised a graphics system for textile design, and worked with graphic methods for urban planning [ny04, ny05".

Palo Alto, California (Established 1964)

The Palo Alto Scientific Center, located adjacent to Stanford University, is one of the original Centers formed in 1964. Originally PASC was headed by Mr. John W. Luke. It has been under the leadership of Dr. Horace P. Flatt since the end of 1968. IBM Fellow Dr. Harwood G. Kolsky was associated with PASC since its begining, promoting projects in physics, microcode, image processing and FORTRAN. In 1990, Randolph G. Scarborough was named an IBM Fellow for his work with FORTRAN.

PASC began its history with three major projects: laboratory automation, nuclear power research, and applied physics in materials science and atmospheric physics. The many reports and papers published by Center staff members since then reflect both the changing times and the influence of the projects which were transfered to Palo Alto from Washington in 1969, and Houston in 1974 when those Centers closed. Project areas that have been significant through the years are listed below. The goal has been to maintain a balance between application and system oriented projects.

Applications of large scale computing: Originally the emphasis at PASC was in high energy, reactor, and atmospheric physics as well as such disciplines as aeronautical, chemical and power engineering. Subsequent studies have included plasma physics, geophysics, the development of a coal gasification model, radiation scattering models for work in solar energy, graphics techniques for use in engineering problems in the public utility industry, and work in wave equation migration anlysis [pa01, pa02". Large commercial data processing problems have also been addressed [pa03".

Whether studying small scale oscillations in a plasma or large scale movements in the atmosphere, the effectiveness of the computer as a tool for experimentation or simulation has been investigated. Research has extended over a range from the basic physical processes, through model formulation, numerical analysis, programming and graphical display techniques to the verification of models by means of measured data.

The Palo Alto center is the only IBM organization to maintain work and interests in large scale computer applications over such an extended period of time. Results have appeared as system studies, application algorithms and in the design of the 3090 Vector Facility. Particularly notable contributions were the Fast Fourier Transform algorithms for solving partial differential equations, now used widely in the seismic analysis area. A number of innovative algorithms have been developed for vector and parallel machine architectures, including A.A. Dubrulle's excellent work on matrix algorithms [pa04", and J. Gazdag and H.H. Wang's cascading method described in a paper in the special Scientific Center issue of the IBM Systems Journal [pa05".

FORTRAN compiler technology: PASC has been the focal point for IBM advances in this technology over the past decade. The vectorizing FORTRAN compiler was especially influential, while the parallel FORTRAN prototype has heavily influenced external standards.

PASC has always realized that the "bread-and-butter" language of engineering and scientific computing continues to be FORTRAN. Some of the Center's best work has come from its support of IBM's customers in this area. Starting in 1976 Randy Scarborough initiated a project, called "FORTRAN Q," to produce a new optimizer. Later improvements were made to the FORTRAN library as well [pa06". The results were made available in September 1978 as a special product [pa07". It then became the basis for the VS FORTRAN product.

There have been so many FORTRAN activities at PASC, but special mention must be made of Alan Karp's performance work [pa08", and Len Lyon's study of data striping [pa09".

Vectorizing FORTRAN Compiler: Probably one of the most fascinating stories in the history of the Scientific Centers is that of Randy Scarborough's vectorizing FORTRAN compiler. This is a classic case in which a "counter-strategic" center project proved to be superior and eventually was accepted as the strategic IBM product.

The project had its roots in a joint study between PASC and Rice University starting in 1981, and finally resulted in the official adoption of the PASC Vectorizing compiler for the new IBM 3090 Vector Processor in 1985 [pa10, pa11". Before the hardware existed, the vectorizing compiler was successfully checked out on N.S. Gussin's 3090 Vector Simulator.

A successful Parallel FORTRAN Prototype which works on multiple 3090's is the most recent extension of this work.

Microprogramming: Successful work on implementing large system functions using writable control store was first shown within IBM at PASC. It led to the world's first high-level-language machine for A Programming Language (APL), PASC products were APL microcode assists for specific models of the System/370, i.e. 135, 145, 138, and 148 [pa12, pa13". The revised Palo Alto design was later the basis of VSAPL, a long-lasting IBM product [pa14". There were several other APL products, listed in the bibliography [pa15, pa16".

The PASC work in microprogramming demonstrated the feasibility of producing large bodies of microcode to perform specific functions. This example encouraged IBM Endicott to develop microcode performance assists for Virtual Machine Facility/370 (VM/370) and its Conversational Monitor System (CMS).

SCAMP - the PC prototype: An outgrowth of work on microprogramming was SCAMP (Scientific Center APL Machine Prototype), the prototype for the IBM 5100 portable desktop computer and a clear predecessor of IBM's later PCs. SCAMP is now in the collection of the Smithsonian Institution [pa17".

Expert Systems Environment: PASC was responsible for the introduction into IBM of expert systems technology and the products built around ESE (Expert Systems Environment). ESE added ease-of-use and new implementation techniques to products based technically on Stanford University's EMYCIN. Products included the IBM ESDE/ESCE Expert Systems.

AIX/370 distributed operating system: PASC introduced IBM's first truly distributed operating system, Advanced Interactive Executive/370 (AIX/370). The project built on earlier university work on the LOCUS distributed operating system, and resulted in the introduction of that technology to IBM as the basis for a new system exploiting IBM's XA architecture.

Digital image processing: Since 1979 PASC has had a series of image-based applications and joint studies, each application being associated with an end-user. The joint studies have resulted in applications such as the new DIMAPS imaging system. Extensive early use was also made of the IBM 7350 Image Processing System (Hacienda). One major study was with NASA concerning the processing of Landsat images [pa18".

Theoretical imaging studies included numerical precision, research in classification techniques, and medical images [pa19, pa20". A successful project demonstrated the practicality of building an image processing system, PCIPS, based on the IBM Personal Computer [pa21". The computer analysis of the Mona Lisa painting has attracted world-wide attention [pa22".

Laboratory automation: An early project at PASC involved the problems of on-line data acquisition. Working with physicists from Stanford University and the Stanford Linear Accelerator Center, PASC staff installed computers that monitored high energy physics experiments and provided immediate graphic output of the results of experiments.

Graphics Information Systems: This influential Palo Alto development effort resulted in four products and defined new applications areas in geographic information systems in public utilities, in manufacturing, and process industries. Four major products have been released: Interactive Geo-Facilities Graphics Support, Graphics Program Generator, Graphics One-Line Diagram, and Composed Document Viewing Utility [pa23".

Networks and software: One programming project involved the rework of the APL Departmental Reporting System and led to a huge number of installations [pa24". The work on PROFS in note-handling and in cross-system communications are currently setting product directions [pa25".

The early work in local area networking led to IBM's first Ethernet(TM) local-area network and contributed significaltly to the technology used in PCnet. The recently announced PSInet computer conferencing system is proving to be valuable in allowing widely-dispersed people to communicate inexpensively in nationwide tests, both in equipment costs and communications call-up charges [pa26".

The Research Support Program: The Research Support Program was started at PASC about six years ago. Now all three US Centers participate in the program and, through it, provide access to their 3090 vector systems to university researchers who have been unable to secure computational support elsewhere. A recent count showed 102 studies active at 63 schools.

Philadelphia, Pennsylvania (1972 - 1974)

In 1972, the IBM New York Scientific Center and the IBM Research Division group on A Programming Language (APL) joined forces and started the IBM Philadelphia Scientific Center. The APL group dated from the early 1960's at Yorktown. The Philadelphia center was headed by Mr. Adin Falkoff, and included IBM Fellow Dr. Kenneth Iverson, the inventor of APL.

The Philadelphia center was closed in 1974 as part of an general economy move. Some of the staff went back to Yorktown, and some went to the APL program development group in Palo Alto.

APL was the major project area for the Center, along with the project area of Management Science taken over from New York. The work on APL resulted in an APL business planning system, extensions to APL itself, and the groundwork in generalized array theory. The APLNET network provided a successful long-standing APL computer center and time-sharing service for other IBM groups.

Management science: Philadelphia assumed many of the management science projects from New York. The group produced two major products, TPACS and EPLAN.

Airline Crew Scheduling - TPACS (Trip Pairing ACS): Jerry Rubin succeeded in producing what is still the state-of-the- art technology for airline crew scheduling, based on careful heuristics and local integer programming. TPACS was sold to the 9 major US Airlines.

Econometric Planning - EPLAN: Taking advantage of the built -in power of APL, Franz Schober designed and implemented a full econometric system with all of the functional capabilities of much larger systems. EPLAN was sold in hundreds of copies worldwide. The econometric group at IBM Corporate headquarters took over the product and succeeded in making follow-on versions of it a standard for important international efforts, such as Project Link with Nobel Prize Laureate L. Klein. EPLAN is still used today at major banking installations [ph01".

There was basic work in mathematical programming, such as on Multi-Commodity Flow Problems by Bill White, and Logical Methods for Integer Programming by Guignard and Spielberg. Other projects were in Information Systems (APLIS), APL business planning system (APLAN), manpower scheduling, generalized inverses, earth resources management, and life insurance budget planning [ph02, ph03".

A mathematical programming system in PL/I and a general APL to PL/I translator were not completed because of the closing of the center.

APL System design and implementation: This was the Center's major project area. This involved experimental implementation of new APL language and system features, such as shared variables for APL (APLSV). Philadelphia also ran a very successful APL computer center that provided time-shared service to many IBM groups, including the other Scientific Centers. This service was a convincing demonstration to IBM management of the value of a well-run interactive time-sharing system to an organization.

Array Theory: In this center Dr. Trenchard More started his long-term research in generalized array theory which was transferred to Cambridge in 1974 and became the basis for A Programming Language 2 (APL2).

Washington (Wheaton, Maryland) (1967 - 1969)

A Center was established July 1966 in Wheaton, Maryland, to be known as the IBM Washington Scientific Center. An earlier Washington Scientific Center had been started briefly in 1964. It was restarted successfully in 1966 with Dr. Horace P. Flatt as manager, only to be closed again in 1969 because resources were not generally available to build a viable center. Most of the Center's staff and work were then moved to Palo Alto.

Significant areas for the Washington Center were environmental sciences, text processing, cartography, and microcode. The environmental sciences area involved work on mathematical techniques for descriptions and mapping of nature, such as weather prediction and water resources - reservoirs, pollution, and waste treatment automation.

★ ★ ★ ★

Summary Comments on the U.S. Centers

In the above sections the reader can find numerous examples of Center work over the past quarter century which anticipated and addressed major IBM and customer problems at just the right time. A common mission allowed the Centers relative freedom within IBM to explore new technical areas, take a broad view of research and applications in science and technology, and make significant contributions in areas of growing importance today.

Retrospective NOTE: In 1989 the authors concluded that much of the strength and productivity of the Centers came from having a Corporate mission, a common Director, local autonomy and long-term stability in their local management. Within a Corporation that is renowned for its change, this has resulted in a staff continuity and technical depth that make the Cambridge, Palo Alto, and Los Angeles Centers well prepared to face the change and opportunities presented to themselves and IBM over their first twenty-five years. Over the next three years, however, IBM faced the severest of strains both within the marketplace and in its financial performance. As we shall comment further, these unprecedented challenges combined with environmental and organizational changes impinging upon the United States Centers led to the situation where in less than three years from the 1989 anniversary -- and the original writing of the above section -- all three Centers found themselves closed. The reader is referred to the section: ***Afterword*** for a discussion of these events.

The World Trade Scientific Centers

World Trade Scientific Centers

The IBM World Trade Scientific Centers had their roots in earlier organizations, as did the Centers in the United States. In 1960 a small group of scientists in Paris, headed by Dr. Rene Moreau, began projects in scientific development for IBM France. In 1965 a small group of scientists from IBM Italy started working at the University of Pisa. However, it was not until 1967, when the IBM Grenoble Scientific Center opened, that Centers based on the model used in the United States made their appearance.

Unlike the US Centers which report to a single director, the World Trade Centers report directly within their respective country organizations. For Europe and the Middle East there has been a Paris headquarters staff organization to help coordinate and support the Centers. Those who have headed this staff over the years have included Conrad Maheu, Ulf Berg, Gary S. Kozak, Raphael Aguilar, Herbert F. Budd, Vittore Casarosa, and Etienne Paris. The present manager is Rodger Hake.

From 1967 to 1972, twelve new Centers were opened around the world. Around 1980, six more were established, and several earlier ones merged with others and were closed. A number of organizations were also derived from the Centers, such as an IBM-sponsored center at the Asian Institute of Technology in Thailand; the European Center for Scientific and Engineering Computing in Rome, Italy; the Numerically Intensive Computation centers in the United States, Europe, and Japan; and the European Networking Center in Germany. While they are different organizations, they carry on the same traditions of scientific exploration and service to IBM and its customers. The following sections can give only a glimpse of the contributions by these dedicated and diverse groups which function uniquely within their own countries.

We call special attention to the Hacienda project as an excellent example of cross-center activity and support involving many Centers. Under the encouragement of IBM Chief Scientist Dr. Lewis M. Branscomb, the project was begun in 1978 with Dr. Patrick E. Mantey of Research San Jose as leader. Hacienda was approved and funded in February 1979 by IBM's Headquarters in Paris.

The project was an experimental, interactive, high- resolution (1024x1024 picture elements) color display with 8 MBytes of storage and associated software [pi03]. It led the way for IBM in image processing and resulted in the IBM 7350 Image Processing System.

The first hardware prototype and the Basic User Software (HBUS) was installed at the Paris Scientific Center in 1981. The next units were shipped one per month to Rome, Palo Alto, Madrid, etc. A product based on the IBM 7350 was announced one year later with HBUS and the High Level Image Processing System (HLIPS) software. The system is still in world-wide use today. Several of the individual Center descriptions which follow comment specifically on this effort.

Bari, Italy (1969 - 1979)

The IBM Bari Scientific Center was established in 1969 with Aldo Canfora as manager. Later it was managed by Attilio Stajano, then Vittore Casarosa. Always a small center, Bari was consolidated with the Rome Scientific Center in 1979.

The Bari center first studied the fields of computer aided instruction and natural language processing. Later it concentrated on relational databases and query languages [ba01].

The major achievement of the center was the development of the A Query Language (AQL) interpreter as a relational database extension to the A Programming Language (APL) systems [ba02].

Bergen, Norway (Established 1986)

The Bergen Scientific Center is one of the most recent Centers to be founded by IBM. It was officially opened on August 26, 1986, with Dr. P.W. Gaffney as manager. The scientists and administration of the Center are housed in the new Information Technology building of the Bergen High Technology Center while the mainframe computing facilities are located in the science building owned by the University of Bergen. The major technologies which are under study at the Bergen center are: vector multiprocessing, parallel processing for the IBM Advanced Interactive Executive (AIX), workstation studies, mapping and technical documentation, and visualization and animation.

The Bergen center's focus has been on project areas of importance to Norway, such as information technology, offshore technology, and oceanographic modelling. In August 1989 the mission of the Center was expanded to cover Environmental Sciences on behalf of the IBM company in Europe. This means that the Center will be the focal point in IBM for environmental modelling and it will provide an information focal point for IBM environmental activities worldwide. The center has attracted a large number of students who come to use the computing facilities [be1, be2, be3].

Cairo, Egypt (Established 1982)

The IBM Cairo Scientific Center was formed through an agreement with the Egyptian Academy of Scientific Research and Technology in March 1982. The Center, headed by Dr. Hamed Ellozy, was officially inaugurated in November 1983. In July 1986 it was relocated to an IBM facility. The present manager since June 1985 is Dr. Mohamed Hashish.

The initial projects of the Cairo center were directed towards environmental studies with various governmental institutions. These projects included modeling the development of Egypt's "new lands," liquid natural gas spills in the Suez Canal, and employment in Egypt [co01, co02].

In May 1985, the first of the projects related to Arabic language research was initiated, namely the development of an Arabic spell-checker capable of checking and assistance with Arabic text. The project partner was the Department of Semitic and Oriental Studies of Cairo University. One of the directions is to develop machine assisted translation from English to Arabic and vice versa. Some of the tools to achieve this have been developed (Lexicon, Dictionaries, Morphology, Syntax). Other Arabic related research projects are in the field of speech recognition & synthesis and multi-language OCR [co03, co04].

There is also research going on at Cairo aimed at achieving a better life for disabled people, such as the rehabilitation of people with visual and auditory disabilities [co05].

Caracas, Venezuela (Established 1983)

In June 1983, the IBM Caracas Scientific Center opened its doors, with a focus on remote sensing and image processing. Dr. Manuel Robollo of Madrid was the first manager.

Now the emphasis is on three areas of scientific study: numerically intensive computing, algorithms and data structures, and software engineering. The center is located in the IBM Building on Avenue Ernesto Blohm, Caracas. The present manager is Dr. Juan Rivero.

Scientific and engineering computations: The studies in astronomy and astrophysics have resulting in publications on electron impact excitation and radiative atomic data [cr01, cr02]. Current work also emphasizes computational chemistry.

Algorithms and Data Structures: Studies have resulted in improving time and space efficiency in generalized binary search trees, the transforming of multi-way trees into a practical external data structure, and new methods for collision resolution in hashing [cr03, cr04].

Software engineering and educational software: This work has added a new in situ distributive external sorting algorithm - EXTQUICK - to the tools available to the software engineer [cr05]. Current work includes in FORTRAN programming environment (FPLUS) and a data base design and control environment.

Grenoble, France (1967 - 1973)

The Grenoble Scientific Center was established in September 1967 as a joint venture between IBM France and the University of Grenoble. Headed by Dr. Jean-Jacques Duby, it was the first IBM Scientific Center outside the United States, and one of the very first example of close cooperation between a French university and an industrial enterprise. It was closed six years later in December 1973 and the staff reassigned.

Activities at Grenoble were devoted mainly to the development of software tools and techniques in two major areas, operating systems and compilers [gr1, gr2, gr3].

Operating Systems: Grenoble played a significant role in the development of CP-67/CMS, in close cooperation with the IBM Cambridge Scientific Center. Some topics studied were: Virtual virtual machines, analysis and control of the thrashing phenomenon, micro-coded CP, CMS file editor on a graphical display, and debugging tools for operating systems.

Compilers: In the compiler area, Grenoble researchers developed conversational compiling techniques achieving incremental compilation of each statement. The results of of this work were used in the FORTRAN Interactive Debug and PL/I Incremental Compiler products.

Haifa, Israel (Established 1972)

IBM Israel Science & Technology (IS&T) is a research and advanced development organization which encompasses two major activities: The Haifa Scientific Center (HSC) and the Haifa Research Group (HRG). The Scientific Center was established August 1, 1972, with Dr. Joseph Raviv as manager. It is located on the campus of the Technion, Israel Institute of Technology in Haifa. Projects of special interest to the customers, planning, medical, and agricultural concerns of Israel have been the major theme in the work of the Center.

Medical applications of computers: Since its establishment HSC has been engaged in medical applications. Various projects with partners have made noticeable progress on improved diagnostic capabilities for ultrasound testing using signal and image processing techniques, estimating cardiac output and gas content of blood, and developing a specialized interactive system for diagnosis of endocrine disorders [ha01, ha02].

Regional planning: A regional planning project conducted by the HSC has contributed to regional rural development through a computer aided modeling and planning system based on mathematical models and linear optimization. This work has been used by the Settlement Study Center at Rehovot, Israel, for whom it was originally developed. The system has also been implemented in Spain and Costa Rica [ha03, ha04, ha05].

Data and image compression: A current project is developing compression techniques both for still images (bi-level, greyscale and color) and for motion video (greyscale and color), as well as contributions to basic compression algorithms. Haifa will utilize this technology in a form processing and storage system for the Swiss Population Census to be conducted in 1991 [ha06].

Signal Processing for the Hearing Impaired: The project demonstrates the applicability of advanced signal processing to appreciably decrease the background noise in assisting the hearing impaired - the "cocktail party" noise problem. A multi-microphone adaptive filtering solution has been formulated and was also applied to the speech recognition problem [ha07].

Agriculture: The Haifa Scientific Center has contributed both to the environment and to the application of IBM's technology in agriculture. A most notable contribution was the development of an interactive system for constructing irrigation time tables, which is now used throughout the country and has an important role in the efficiency of agricultural production in their arid lands. This became a Country program offering in March 1979 and was used throughout Israel in the agricultural settlements [ha08, ha09].

Software Systems & Algorithms: One HRG department specializes in programming languages and compilers, distributed processing and operating systems, performance evaluation, and on-line presentation. A notable contribution was the development of a REXX compiler prototype, which is the base for the recently announced VM/REXX compiler. Haifa has also had a leading role in the "Multi-target English Translator" [ha10] and "Optimizing compiler" projects.

Computer Engineering & Signal Processing: HRG projects in this area have focused on speech processing, image processing, VLSI design verification, and testing tools. Verification methodology developed at Haifa has been used by the Advanced Engineering Systems Development group in Austin [ha11, ha12, ha13, ha14]. Physical design tools have been used in Burlington and Austin and E-Beam Large Array Macrotizing tools in East-Fishkill.

Applied mathematics: HRG projects in Applied Mathematics include the theoretical and practical applications of numerical analysis to problems arising in mathematical physics and engineering, such as stiff differential equations and partial differential equations with complex boundary conditions. Work has included the development of algorithms for problems arising arising within IBM development laboratories, such as semiconductor simulation and Nastran problems [ha15, ha16, ha17].

Heidelberg, Germany (Established 1968)

When the IBM Heidelberg Scientific Center (HDSC) first opened in January 1968, its projects concerned the physics of bubble chambers and scintigraphic photo analysis. It was temporarily located in Mannheim before moving to its new building in Heidelberg. Dr. Willy Kattwinkel was the first manager.

Since 1973, when Dr. Albrecht Blaser became manager, the areas of interest are more closely aligned with data processing: capabilities for the processing of written and spoken natural language, knowledge based systems, relational data models and query languages, visual languages for end users, additions to A Programming Language (APL) products, numerically intensive computing (NIC), and immunological information processing for transplant surgery. The Institute for Knowledge Based Systems was added to IBM Germany's Science Center in November 1988.

Natural language database query: From 1974 to 1982 HDSC developed and validated a technology which allows limited natural language interaction for database queries, with support for many western languages. By exchanging the grammar and vocabulary only, one can accommodate German, English, French, etc. [hd4, hd5, hd6].

Non-First-Normal-Form (NF2) Relational Data Model: NF2 extends the classical relational data model by essentially dropping the First Normal Form condition. It allows complex data structures as field values of tables and encompasses the relational and the hierarchical model as special cases.

HDSC implemented a prototype called the Advanced Information Management Prototype (AIM-P), a program to add extended non-first-normal-form (NF2) and other capabilities to the relational data model and to IBM's Structured Query Language (SQL). (hd3, hd7, hd8).

Extensions to Query-By-Example (QBE): Since the mid 1970's, HDSC has had a long-standing research effort to add visual programming for end users to relational database systems. In 1978 Interactive Programming by Endusers (IPE) was an early spread-sheet for developing application programs by designing forms. Extended Query by Example is a prototype language and supporting system for querying data from a relational database. (hd2).

Graphical, direct manipulation languages: More recently, human factors investigations have led to development of a methodology for solving application problems through graphical manipulation. Based on human factors investigations, HDSC has developed a methodology of graphical manipulation for solving application problems. The methodology is being applied in a variety of applications, and has influenced the IBM Common User Interface Architecture [hd9, hd10].

Numerically intensive computing: HDSC has cooperative projects with universities and commercial customers to use NIC for reactive flow; chemical, mechanical, and financial modeling; and simulation, in concert with the European Supercomputing Initiative. [hd11].

Information system for transplant surgery: Jointly with the Institute for Transplant Immunology of the Heidelberg University an interactive, distributed information system was developed for global use to support immunological research in transplant surgery across three hundred transplant centers worldwide, using an interactive and distributed information system [hd12].

APL Products: As a result of first-hand, concentrated use of APL at HDSC, additional capabilities were developed for APL and APL2 in support of complementary functions, interlanguage communication, and graphic presentation. This resulted in products, such as the European Program Offering "APL Complementary Functions", the APL Interlanguage Communication Facility function in the APL2 program product, and the German Program Offering "APL System for Graphic Presentation".

Institute for Knowledge Based Systems: From 1984 to 1987 a knowledge based system (LEX) was implemented capable of using legal knowledge and answering questions on traffic law in German which were also asked in German. (hd13). From 1986 to 1988 the Institute at Heidelberg and Stuttgart developed a prototype for German text understanding. Since 1987 the Institute has taken part in the European Research Initiative EUREKA working together with Swiss and Belgian enterprises and Swiss and German universities [hd1, hd14, hd15].

Kuwait City, Kuwait (Established 1980)

The IBM Kuwait Scientific Center was founded in September 1980 with the arrival of the first Center manager, Dr. Samar Attasi, and two assignees. KSC operates in cooperation with the Kuwait Institute for Scientific Research (KISR). The Center is located within the premises of KISR, and is its partner in many major projects. The Center supports GULFNET, the academic network in the Gulf Area. The present manager is Dr. Sherif S. El-Dabi.

The first projects at KSC were on image processing and environmental studies. KSC had early use of an IBM 7350 (Hacienda), which made high quality satellite remote sensing studies possible. Arabic studies were added in 1983. The present project areas are:

Arabic language related studies: Under Arabic studies there have been projects in speech synthesis and recognition, computational linguistics and databases, a course authoring system (DA_RES), a Script I/O system, printed character recognition, Arabic advanced function printing, and desk top publishing [ku01, ku02, ku03].

An IBM Personal Computer-based Arabic speech synthesizer has been designed that is capable of synthesizing unlimited vocabulary or Arabic text. The system is based on two major components, a speech synthesis unit for Arabic language, and complete text-to-speech rules.

A comprehensive Arabic morphological analyzer/generator system with a root-based Arabic lexical database was created for verification and support of other linguistic components. It will be integrated into text retrieval packages (STAIRS) to add to it morphological search capabilities.

Image processing applications: Projects have been in satellite remote sensing, medical and PC-based applications. KSC is the only group in Kuwait that has equipment dedicated to remote sensing applications. An ongoing major remote sensing project has been with the survey and mapping of Kuwait using the most recent Landsat 5 satellite data. It included studies of surface features, marine sedimentation, underground water, and the desertification problem of Sudan's Red Sea hills [ku04, ku05].

A thermal infrared image processing system was developed. It was used for a joint project with Kuwait University to study oil distillation columns. Another modeling project concerned the characterization of micro pores in catalysts.

Environmental sciences: Work has been completed on passive cooling and air pollution The focus is now on oil pollution and source identification. One activity involves oil spill identification using laser fluorescence technology. Another project is concerned with the data analysis of a reverse osmosis desalinization plant [ku06, ku07, ku08, ku09].

Madrid, Spain (Established 1972)

The Autonomous University of Madrid is host to the IBM Madrid Scientific Center, which was chartered in 1972 through an agreement between IBM Spain and the University. The University furnishes premises for the center and IBM cooperates with computer facilities for research projects. The first manager of the Center was Dr. Rafael Aguilar. He was succeeded by Dr.J.L. Picon and Dr. Victor Martinez (who is also the present manager), and Dr. Jose Louis Becerril. The main products of the Madrid Center have been:

APL interpreters and interpreter writing technology: Madrid developed the very popular IBM PC APL (versions 1 and 2). It also developed the 5550 Nihongo (Japanese) APL (versions 1 and 2), the JX Nihongo APL, and APL2 for the IBM PC (two versions). Next to the Hacienda project, Madrid's contributions to the Japanese PC software effort is probably the best example of international collaboration in the Scientific Centers [ma1, ma2, ma3].

Image processing: Madrid was a leader in the development of numerous applications of image processing in Europe. Among its study partners at the Autonomous University was the Nobel Laureate Professor S. Ochoa, head ot the Center of Molecular Biology.

Madrid was one of the Center partners in the Hacienda (IBM 7350) project, and in the production of the HLIPS (High Level Information Processing System) software [ma4, ma5, ma6].

Natural language projects: The center developed the Spanish and Catalan dictionaries for text processing products. They made linguistic functions available in the current line and continue the development of new functions [ma7, ma8].

PC Bank Printers and Networks: The Printer Banking Adapter developed at Madrid was announced by Hursley in 1984. The Financial B-loop Adapter was announced by Boeblingen in 1986.

The Madrid Center has also done many applications in other fields such as operations research, automated greenhouses, bacteriology, and fluid dynamics [ma13, ma11, ma12, ma13].

Mexico City, Mexico (Established 1971)

The Mexico Scientific Center has its origins in 1971, when it was called the "Latin American Scientific Center," with James S. Nist as its first manager. The center really became a Scientific Center under the leadership of Dr. Baxter H. Armstrong, who was manager from 1973 to 1975. Its name was changed to the IBM Mexico Scientific Center, in keeping with its close ties to Mexican projects.

For many years the Center was located near the main campus of the National University in southern Mexico City. The present location, inaugurated in 1986, is a larger, fully equipped building two blocks from the IBM Mexico headquarters building near downtown Mexico City. The present manager is Dr. Eduardo de la Paz.

Mexico's Scientific Center mission is to establish projects of scientific and technological research with government agencies, universities and research institutions in order to collaborate to solve priority problems of the country. Over the years the Center has gained a strong reputation within Mexico in several areas of specialization.

Major project areas: These include artificial intelligence and expert systems, image processing and remote sensing, geographical information systems, statistics and applied mathematics, and data base design. Of special interest is the work on geographical information systems that support decision making concerning natural resources by allowing the mixing of different kinds of mapped data, such as temperature, rainfall, and coastal dune vegetation [me1, me2, me3].

Geographical information system: A system which can manage different maps of information (temperature, rain, natural resources, etc). It has been used as a support for those institutions which are involved in decision making for the natural resources administration.

Application software packages: A number of programs have also been built for the IBM Personal Computer (PC) and Personal System/2 (PS/2) for time series analysis, satellite data processing, and database consistency checking.

SDE & SDF is a PC program designed to help economists and investors perform data analysis through time series and graphs. SPIPR is a powerful software system that allows processing of satellite images on a PS/2. CEST is an expert system in time series analysis, and DETEC is a software package designed to run on a PS/2 and a mainframe to allow one to detect inconsistencies in a knowledge base.

Paris, France (Established 1977)

The present IBM Paris Scientific Center was officially chartered in 1977, but it built on a tradition of projects that began in January 1960 with the IBM France Scientific Development Department, headed by Dr. Rene Moreau and followed by Alain Croissier. The present manager is Dr. Claude Hans.

The main areas of exploration since the Center was founded have been mathematics, linguistics, image processing, and artificial intelligence. In the early years, Paris pioneered in mathematical programming and non-numerical information processing, such as computational linguistics and computer algebra [pr1, pr2, pr3]. The main research domains over 30 years have been:

Mathematics: Important contributions were made in the 1970's in the area of qualitative statistics, for the analysis of qualitative variables. More recent developments have shed a new light on the problem of lengthy complex computations precision. [pr4, pr5].

Linguistics: Following the early work in statistical linguistics, many Paris projects resulted in operational software. For instance, the Universal Language Generator, an IBM internal use program, was mainly developed in the 70's at Paris (it was later transferred to the Heidelberg Scientific Center). The Deaf Children project, using a PC equipped with a special voice processing card to help speech train young children with serious hearing impairment resulted in the IBM Speech Viewer product, was announced in 1988 for the PS/2. The Stenotypy project, aimed at automatic translation of stenotyped text into French natural language, produced TASF (Traduction Automatique de Stenotypes en Francais) a Country Product Offering.

The French thesaurus project developed lexicographical techniques used for multiple native languages features of the DW370 and Office Vision Advanced Text Processing products from the IBM Bethesda Laboratory, together with the dictionary and syntax data pertaining to the French language. Still at the research stage is the Automatic Dictation project of the Paris Scientific center, using as a basis the Information Theoretical methodology and the Voice Processing hardware developed in Yorktown Heights.

Image Processing: Paris was the first Center partner to receive the hardware prototype Hacienda (IBM 7350) and the Basic User Software (HBUS), in 1981. The projects at Paris resulted in the High Level Image Processing System (HLIPS) and APL Image Processing Attachment Support (APLIAS), which both became products.

In application joint studies with outside partners, the Center studied different techniques for satellite imaging, some of which were later used in commercial satellite support systems (e.g. SPOT) [pr6, pr7].

Artificial Intelligence: Paris was an early pioneer in this field, with one of the first Expert Systems developed in the 1970's, "Bateau sans medecin" (Ship without a medic). On the product side, Programming in Logic (PROLOG) for VM and MVS was developed in Paris for using logic programming on IBM machines. The VM-PROLOG Product Offering was developed in Paris. MVS-PROLOG is derived from it. More recently announced, the "IBM PROLOG for 370" Program Product was based on a totally new prototype developed in Paris and transfered to the IBM Menlo Park Laboratory.

More recent work focuses on theoretical breakthroughs in non-classical logic through development of a qualitative logic model allowing Order of Magnitude Reasoning; the model was validated using non-standard analysis and applied to one of the first expert system based on non-shallow reasoning [pr8, pr9].

Pisa, Italy (Established 1971)

In 1965 the Centro Studi IBM was established in Pisa, with Dr. Renzo Marconi as manager. Its mission was to cooperate with the University of Pisa in creating and managing the Centro Nazionale Universitario di Calcolo Elettronico (CNUCE), a large Computing Center encouraging data processing in the Italian academic and research environment.

In 1971 the Centro Studi became the Pisa Scientific Center and moved to its present building on via S. Maria, the street of Galileo's old study. The center has major project activity in surface hydrology, computer networks, and econometric software. The emphasis on surface hydrology was a result of the very serious flood of the Arno River which occurred in 1966.

Surface hydrology: The model of the river Arno was the final result of the Surface Hydrology project for the Tuscany Regional Authorities. It was a large program for simulating the rainfall runoff behavior of river basins, in order to evaluate the effects of expected floods. Experiments were also made with the river Nile and Lake Aswan, in Egypt [pi01].

RPCNET: RPCNET was an effort to develop a network for VM-based computing centers, in cooperation with CNUCE and Academic and Research partners in Bari, Bologna, Padova and Turin. The Pisa center played a leading role in promoting, designing and implementing the RPCNET Network. The network has been running successfully for several years, linking the computing centers of all the parties.

Econometrics: Pisa designed and implemented a complete set of programs to estimate, validate, analyze, and simulate both linear and nonlinear reliable macro- econometric models. The software has been made available for research purposes to several scientific organizations all around the world [pi02].

IBM 7350 (Hacienda): The special contributions of the Pisa center to the 7350 project were the design and implementation of HBUS (Host Basic User Subroutines), which is the host-resident subroutine library for the IBM 7350 Image Processing System [pi03].

Medical imaging: The expertise in Image Processing developed on the IBM 7350 effort was the basis for a cooperative project with the University of Pisa (Project Heart), to analyze radiographs related to the internal innervation and blood circulation of the cardiac muscle. Experiments in 2D and 3D object detection and recognition were completed with the University of Florence [pi04].

Volcanic risk modeling: Image Processing and Mathematical Modeling were used for a research project on volcanic phenomena. In cooperation with the University of Pisa, a project for automatic drawing of Volcanic risk maps, related to volcanic eruptions was completed. Experiments have been made, to validate the mathematical methods used by the project, with historic eruptions of Vesuvius and Etna in Italy, and Mount St.Helens [pi05, pi06].

Italian text processing: A major result has been the development in 1988 of a Country Program Offering for Italy, named Sistema C, a VM system which performs interactive sophisticated checking and correction of Italian texts in the XEDIT environment. The analyses made by Sistema C relies on a large computerized dictionary, which includes several functions related to the grammatical classification of Italian words [pi07].

Computer Aided Translation from English to Italian, is another project in its early stages. The goal is human-aided machine translation of English into Italian, with special attention is being given to the translation of IBM Manuals.

Brazil SC, Brasilia (1980 - 1986)

The IBM Brasilia Scientific Center was established in 1980 in the new capital city of Brazil, with Dr. Jean-Paul Jacob as manager. It's focus was on enhancing the quality of life, the work environment, and the overall community's well-being. In 1984, the Center was honored for this work by receiving the American Chamber of Commerce Award for Corporate Service to the Community. In 1986, the Center merged with other activities in Brazil and moved to Rio de Janeiro.

Brazil SC, Rio de Janeiro (Established 1986)

The IBM Rio Scientific Center was established in 1986 when the Center in Brasilia moved to Rio de Janeiro. All research and educational activities of the Latin American Systems Research Institute (founded 1980), the Software Technology Center (founded 1982), and the IBM Brazil Sumare Plant were integrated into the Rio Scientific Center under the present manager Fernando Borges.

Since the consolidation of the activities in Rio, priorities are now more business-oriented, and current research areas are databases, data communications, artificial intelligence, logic programming, microelectronics, software engineering, vector processing, and advanced signal processing.

The Center conducts educational activities in the areas of computer systems, software engineering, manufacturing technology and quality.

Significant social and economic contributions have been made through development of an accurate model of the borer insect and sugarcane ecosystem for plague control, and through study by remote sensing of the environmental impact of the Tucurui hydroelectric facility.

Contributions to computer science include the PACCHIP system, an integrated computer-aided design (CAD) package for very large scale integrated (VLSI) chip design, the CHRIS software tool for database conceptual design and prototyping, and an expert system for heuristic learning of medical diagnoses [ri01, ri02, ri03].

Rome, Italy (Established 1979)

The IBM Rome Scientific Center was chartered in 1979 under the leadership of Pier-Luigi Ridolfi. It absorbed the projects and staffs of the Centers at Bari and Venice and some personnel from Pisa. The present manager is Dr. Paolo Franchi.

At first the projects of the Rome center were a continuation of those from the former centers. Then over time the project focus shifted to new areas as, as described below. For example, relational database and query languages continued in Rome in the form of territorial data base, and the activity on hydrology and air pollution became more general and addressed other aspects than those pertaining to the Venice area [ro01, ro02, ro03, ro04].

Distributed processing: Based on the experience gained from the Pisa RPCNET project, the Rome Center started a new project, VMDF (Virtual Machine Distributed Facility). VMDF was used on XCF (eXtended Communication Facility) to create a single-image VM system that allowed file sharing between remote virtual machines and catalog sharing between VM control programs. It was operational on the network for the Centers in Italy until 1984 [ro05, ro06].

Image processing: The Rome center participated in implementing the IBM 7350 Image Processing System, mentioned in the introduction, and since 1981, has used its special color capabilities, along with the IBM 3838 Array Processor, in numerical modeling projects such as seismic wave migration, fluid dynamics, and meteorology [ro07, ro08, ro09, ro10, ro11, ro12].

Natural Language processing: Natural language processing is among the new research directions with special relevnce to Italy. One project began in 1984 with the goal of understanding the Italian language, and in 1987 resulted in the most complete text understanding system presently available in Italy [ro13, ro14].

Speech: A text-to-speech synthesizer project for Italian, was begun in 1982 using IBM hardware. In 1984 a speech recognition project for Italian was started with the help of the Yorktown Continuous Speech group. By early 1988 the Rome Scientific Center, working on the peculiarities of the Italian language, reached the goal of demonstrating a large vocabulary (20,000 words) recognizer for Italian, a unique achievement in terms of performance, for a language other than English [ro15, ro16].

ECSEC: The skill in numerically intensive computing (NIC) such as water system modeling, air pollution applications, seismic modeling on IBM 3838, fluid dynamics computational physics, etc., developed over more than a decade in the Italian Scientific Centers, resulted in the establishment in 1984 of the European Center for Scientific and Engineering (ECSEC) as a department of the Rome Center. Since then the Center evolved to become the European focal point in IBM for NIC along with Kingston, PASC and Tokyo NIC Centers.

Tokyo, Japan (Established 1970)

The Tokyo Scientific Center was founded January 1, 1970, with E. Konno as its first manager. In January 1983, the Center became a part of Japan Science Institute, now called the Tokyo Research Laboratory. Toyohisa Kaneko is the present manager. Main project areas have been natural language support, image processing, graphics and displays, scientific and engineering computation, and education as described below.

Natural Language: The natural language studies have included research in machine translation between English and Japanese. In the late 1970's there was also work done on converting Kana to Kanji for a Japanese word processor [to01, to02].

Image Processing: The Tokyo Center has done many applications of image processing as joint studies with government and university groups emphasizing regional and coastal planning [to03, to04]. More recent studies have included optical character recognition (both printed and handwritten), and color image databases.

Graphics and Displays: Projects in graphics and displays have included the development of the application programming interface (API) for the IBM 5080 display and its follow-on, and the development of applications for rendering, particularly texture mapping.

Scientific and engineering computation and visualization: Projects in this area have included chemical computer-aided design, computational chemistry, and engineering simulation, such as air flow simulation in a clean room. New graphics techniques have been applied to the rendering of molecules [to05, to06].

Education: Since its beginning the Tokyo Center has always conducted projects with university partners. Recent projects have concerned assistance for the handicapped, mathematical computer-aided instruction, and campus networking [to07].

UKSC, Peterlee, England (1969 - 1979)

The IBM United Kingdom Scientific Center was originally founded in Peterlee, Durham, in May 1969, under the leadership of Dr. Colin J. Bell. In 1979 the Center altered its focus considerably, and was moved to Winchester. The major project areas at Peterlee were relational databases and local and regional planning.

Relational data base: The relational database work culminated in the early and successful "Peterlee Model". This data base program was the most famous of the Center's products. It was IBM's first successful relational data base program used by customers and did much to popularize the concept [pe01, pe02].

Local and regional planning: For ten years the Center had many projects with local and regional government groups on problems involving the practical use of computers for planning in their organizations.

UKSC, Winchester, England (Established 1979)

In 1979 the IBM United Kingdom Scientific Center (UKSC) adopted a new charter and a new project focus, and moved from Peterlee to Winchester. That new charter established a new set of project areas for the Center: image processing, graphics systems, and speech processing. The major contributions in these areas are outlined below. In 1988, the Center was a joint winner of the British Computer Society Award for applications.

UKSC is unique in that it has a formal connection to an IBM laboratory - IBM United Kingdom Laboratories Limited, at Hursley Park, England. It is the only Center with such a formal connection. UKSC also is unusual in its extensive use of visiting professionals, forming roughly half of the staff. The visitors work at UKSC for periods between one and three years.

The UKSC Winchester managers have been Peter D. Atkinson, Goeffrey W. Robinson, Ian F. Brackenbury, and Rodger B. Hake. Dave Steventon is the current manager.

Image processing: This area grew out of early work at UKSC in molecular graphics. As the equipment improved the emphasis has been particularly the high content, high resolution end of the business as used in medicine and satellite imaging. UKSC has concentrated on the application of its high-content image work to medical applications of imaging the heart and neurons. This led to the Interactive Applications eXecutive (IAX) Program offering which is 370 VM/CMS-based [uk01, uk02].

Graphics: The focus in graphics has been on modeling using construction solid geometry. The WINSOM system has been developed to make realistic pictures on a raster display terminal. The input to the program a series of building blocks of single regular primitives such as cubes, spheres, and cylinders. These are combined to produce good quality, rendered images of complex objects [uk03].

Speech synthesis and natural language processing: The concentration at the UKSC been on text to spoken English synthesis, and it has established what is believed to be the best center of expertise in the Company. The automatic phonetic transcription from input text has been particularly studied [uk04, uk05].

Parallel Processing: UKSC has been studying the generation of WINSOM (see above) images at high speed using arrays of transputers.

Pharmaceutical projects: A number of projects have focused on the use of graphics for molecular modelling and related topics, now of considerable interest to the pharmaceutical world. The UKSC also is working on the use of expert systems in clinical trials.

Scientific and Technical Data Visualization: The UKSC's graphics and imaging expertise has been exploited in number of fields involving the use of graphical techniques for the presentation of forms and the visualization of large amounts of complex data. Areas for study have been solid geometry, raster display of complex three-dimensional constructions, molecular modeling for the pharmaceutical industry, high speed imaging using arrays of transputers, automatic interpretation of images, stereoscopic computer vision, fluid flow, astronomy, archaeology, theoretical physics, molecular beam epitaxy, and liquid crystal structure. More recently the techniques are also being used for computer generated art [uk06, uk07].

Computer Vision: Imaging expertise has and is being applied to the automatic interpretation of images. The UKSC has made a number of advances in the area of stereoscopic computer vision.

Venice, Italy (1969 - 1979)

The IBM Venice Scientific Center was located in a thirteenth century palace - the only Center to be so honored. The center was established in 1969 with Luciano Lippi as manager.

As was the case in Pisa, a natural disaster had an important influence in the founding of the Venice center. A very serious storm in 1966 caused a storm surge of 194 cm in Venice and drew world-wide attention to the problems of saving Venice. This prompted IBM Italy to join in the government-sponsored effort by establishing a scientific center there.

As a result, much of its work centered on hydrology and air pollution. The numerical modeling of combined tidal and meteorological effects in the Venice lagoon, and of the sinking phenomenon in the Venice area due to the pumping of groundwater, remain as the major contributions of this Center to basic knowledge in these fields [ve1, ve2].

These results have been widely published in articles and books and are still used as a term of reference by the scientific community investigating hydrological aspects of the environment. The Venice Center was closed in 1979 and became part of the new IBM Rome Scientific Center.

Selected References

Selected References

Bari, Italy (1969 - 1979)

ba01 F. Antonacci, P. Dell'Orco, V. N. Spadavecchia. "AQL: an APL-based System for Accessing and Manipulating Data in a Relational Data Base System," Proc. of APL 76 Conference. Ottawa, Sep 1976. Ed.: G. Truman Hunter.

ba02 P. Dell'Orco, V. N. Spadavecchia, M. King. "Using Knowledge of a Data Base World in Interpreting Natural Language Queries." Information Processing 77 (Proc. of IFIP Congress 1977). Toronto, Aug 1977. Ed.: B. Gilchrist. North Holland Publ. Co.

Bergen, Norway (Established 1986)

be01 A.D. Jenkins, J. H. Salvesen, "Ocean Current Modelling at the Bergen Scientific Centre," ODAP Current Modelling Workshop, Roros, 1988. Trondheim: Oceanor - MOMOP Report No. 1.

be02 A. Kamel, "Elastic Modeling on the IBM 3090 Vector Multiprocessor," European Conference on Mathematics in Industry (ECMI), Glasgow, 28-31 Aug 1988.

be03 L. de Pillis Lindheim, J. Petersen, J. de Pillis, "An Iterative Solution to Special Linear Systems on a Vector Hypercube," 3rd Conf. on Hypercube Concurrent Computer and Applications, Vol II, 1988, ACM Press.

Cairo, Egypt (Established 1982)

co01 M.A. Serag El Din, "Spread of NG Following a Small LNG Spill from a ship-Tanker," Proc. of the 4th Int. Conf. on Num. Methods in Thermal Problems, Swansea, U.K., Jul 1985.

co02 Hisham El-Shishiny and Bayoumi Attia, "Multi-Objective Modeling for the Planning and Management of New Lands in Egypt - A Case Study," IFAC Conference on systems analysis Applied to Water and Related Land Resources, Lisbon, Oct 1985.

co03 O.S. Emam and M.A. Hashish, " Application of Hidden Markov Models to the Recognition of Isolated Arabic Words," 10th National Computer Conference and Exhibition, Jeddah, Saudi Arabia, Feb 28 - Mar 2, 1986.

co04 K. Abdel Hamid, M. Hashish, O.Emam, D. Kamal, A. El-Hetw and K. Ahmed, "A Multi-Feature Computer-Based Speech-Training System for Deaf Children," IEEE/EMBS 9th Annual Conf., Boston, Nov 1987.

co05 H.Y. Abdelazim, M.A. Hashish, "Automatic Reading of Bilingual Typewritten Text," CompEURO '89-VLSI and Computer Peripherals IEEE conference, Hamburg, May 8-12, 1989.

co06 O.S. Emam, M.A. Hashish, "Large-Vocabulary Isolated Arabic Word Recognition System: Preliminary Results," EUROSpeech '89, Paris, Sep. 26-28, 1989.

co07 T. El Sadany, M. Hashish, "Arabic Morphological System," (accepted for the IBM Systems Journal, 25th anniversary for Scientific Centers issue).

Cambridge, Massachusetts (Established 1964)

ca01 R.A. Meyer & L.H. Seawright, "A Virtual Machine Time-Sharing System," IBM Systems Journal 9, No. 3, 1970.

ca02 R.P. Parmelee et al. "Virtual Storage and Virtual Machine Concepts, "IBM Systems Journal 11, No. 2, 1972.

ca03 R.J. Creasy, "The Origin of the VM/370 Time-Sharing System," IBM J.of Res. & Develop. 25, No. 5, Sep 1981.

ca04 L.H. Seawright & R.A. MacKinnon, "VM/370 - A Study of Multiplicity and Usefulness," 18, No. 1, 1979.

ca05 Daniel J. Oberst and Sheldon B. Smith, "BITNET: Past, Present, and Future," EDUCOM Bulletin, 21, No.2, Summer 1986.

ca06 E.C. Hendricks & T.C. Hartmann, "Evolution of a Virtual Machine Subsystem," IBM Systems Journal 18, No. 1, 1979.

ca07 L.H. Holley, R.P. Parmelee, C.A. Salisbury, & D.N. Saul, "VM/370 Asymmetric Multiprocessing," IBM Systems Journal 18, No. 1, 1979.

ca08 F.T. Kozuh, D.L. Livingston, & T.C. Spillman, "System/370 Capability in a Desktop Computer," IBM Systems Journal 23, No. 3, 1984.

ca09 Y. Bard, "The VM Performance Planning Facility (VMPPF)," Computer Measurement Group (CMG) Trans. 53, 53-59, Summer 1986.

ca10 H.G. Kolsky and R.A. MacKinnon, "History and Contributions of the IBM Scientific Center," IBM Systems Journal, 28 No.4, 502-524, 1989.

ca11 R.A. MacKinnon, "The Last Report 1991-1992 of the Cambridge Scientific Center," IBM Form Number G325-4101, 1992. This document is available to the public from IBM.

Caracas, Venezuala (Established 1983)

cr01 C. Mendoza, "Electron impact excitation of 2p q and 3p ions," Atomic Data Workshop: Low Energy Collision Theory Techniques for Atomic Excitation and Radiative Data, ed. W.B. Eissner, Proc. DL/SCI/R24 (SERC Daresbury Laboratory, UK), p.80, 1986.

cr02 C. Mendoza and C.J. Zeippen, "Radiative atomic data for neutral magnesium. I. Oscillator strengths," Astron. Astrophys. 179, 339, 1987.

cr03 W. Cunto and J.L. Gascon, "Improving time and space efficiency in generalized binary search trees," Acta Informatica, 24, 583, 1987.

cr04 W. Cunto and P.V. Poblete, "Two Hybrid Methods for Collision Resolution in Open Addressing Hashing," Scandinavian Workshop on Algorithm Theory, Lecture Notes in Computer Science 318, Springer Verlag, p.113, 1988.

cr05 W. Cunto, G.H. Gonnet, J.I. Munro, "EXTQUICK: an in situ distributive external sorting algorithm," Proc. VI Intl. Conference of the Chilean Computer Science Society, p.61, 1986.

Grenoble, France (1967 - 1973)

gr01 J. Bellino, C. Hans, "Virtual machine or virtual operating system?," ACM-SIGARCH-SIGOPS, Harvard, 1973.

gr02 Nguyen Thanh Thi, "A graphic conversational aid for errors searching and teaching of an operating system," ACM-SIGARCH-SIGOPS, Harvard, 1973.

gr03 J. Rodriguez-Rosell, "Experimental data on how program behavior affects the choices of scheduler parameters," ACM, Stanford, 1971.

Haifa, Israel (Established 1972)

ha01 A. Angel, S. Gal, J. Raviv, "Optimal Procedures for Sequential Discrimination between Two Classes," Proc. of MEDINFO 80, Tokyo, Oct 1980.

ha02 D. Chazan, S. Gal, "A Markovian Model for a Central Blood Bank," Management Science, 23, Jan 1977.

ha03 S. Gal, M. Landsberger, B. Levykson, "A Compound Strategy for Search in the Labor Market," Intl. Economic Review, 22, No. 3, 597-609, 1981.

ha04 Y. Ben-Ari, S. Gal, "Optimization of Replacement Policy in a Dairy Herd," European J. of OR, 23, 2, Feb 1986.

ha05 M.C. Golumbic, M. Markovich, U.J. Schild, S. Tsur, "A Knowledge-Based Expert System for Student Advising," IEEE Trans. on Education, E-28, 120-124, Mar 1986.

ha06 E. Walach, J. Bruck, D. Chevion, E. Karnin, D. Ramm, "A Modified Block Truncation Coding Technique for Image Compression," Proc. of Intl. Conference on Advances in Image Processing and Pattern Recognition, Pisa, Italy, Dec 1985.

ha07 D. Chazan, Y. Medan, U. Shvadron, "Noise Cancellation for Hearing Aids," Proc. of ICASSP 86, Tokyo, Japan, Apr 1986.

ha08 D. Ramm, D. Chazan, "A Mixed Numerical-Analytical Method for Underground Water Flow and Tracer Displacement," Water Resources Research, 16, 1980.

ha09 E. Gofman, M. Rodeh, "Incorporating Unknown Pipe Characteristics into Loop Oriented Hydraulic Network Solvers," J. Hydraulics Div., ASCE, 107, No HY9, Sep 1981.

ha10 I. Golan, S. Lappin and M. Rimon, "An Active Bilingual Lexicon for Machine Translation," Proc. of COLING'88, the 12th Intl. Conference on Computational Linguistics, Budapest, Aug 1988.

ha11 I. Berger, Z. Kohavi, "Fault Diagnosis in Combinational Tree Networks," IEEE Trans. of Computers, C-24, No. 12, Dec 1975.

ha12 D. Bernstein, M. Rodeh, I. Gertner, "On the Complexity of Scheduling Problems for Parallel/Pipelined Machines," IEEE Trans on Computers, Dec 1986.

ha13 A. Peled, "Digital Signal Processing Hardware," Invited Review, IEEE National Convention, I., Oct 1977.

ha14 S. Wimer, R.Y. Pinter, J.A. Feldman, "Optimal Chaining of CMOS Transistors in a Functional Cell," IEEE Trans. on CAD, CAD-6, No. 5, 795-801, Sep 1987.

ha15 I. Efrat, M. Tismenetsky, "Parallel Iterative Linear Solvers for Oil Reservoir Models," IBM J. of Res. & Develop. 30, 2, 184-192, Mar 1986.

ha16 S. Gal, B. Bachelis, "An Accurate Elementary Mathematical Library for the IEEE Floating Point Standard," ACM Trans on Mathematical Software, Jun 1987.

ha17 A. Ziv, V. Amdursky, "On the Numerical Solution of Stiff Linear Systems of the Oscillatory Type," SIAM J.Applied Math, 33, No.4, Dec 1977.

Heidelberg, Germany (Established 1968)

hd01 G. Bouma, E. Koenig, H. Uszkoreit, "A Flexible Graph- Unification Formalism and its Application to Natural Language Processing," IBM J. of Res. & Develop., Mar 1988.

hd02 A. Blaser et al, "Integrated Data Analysis and Management System (IDAMS) - Feature Description," in J.W. Schmidt, M.L. Brodie (eds.): Relational Data Base Systems. Springer-Verlag, 1983.

hd03 P. Dadam, R. Erbe, J. Guenauer, V. Lum, P. Pistor, G. Walch, H.-D. Werner, J. Woodfill, "Designing DBMS Support for the Temporal Dimension," Proc. of the 1984 ACM SIGMOD Conference, Jun 18-21, 1984, Boston, Mass. ACM Order No. 472840.

hd04 H. Lehmann, N. Ott, M. Zoeppritz, " User Experiments with Natural Language for Data Base Access," Proc. of the 7th Intl. Conference on Computational Linguistics (COLING'78), Bergen, Norway, Aug 14-16, 1978.

hd05 Hub. Lehmann, N. Ott, M. Zoeppritz, "A Multilingual Interface to Databases," IEEE Database Engineering Bulletin 8, No. 3, Sep 1985.

hd06 M. Zoeppritz, "Human Factors of a Natural Language Enduser System," In A. Blaser, M. Zoeppritz (eds.): Enduser Systems and their Human Factors. Lecture Notes in Computer Science 150, Springer-Verlag, 1983.

hd07 G. Jaeschke, H.-J. Schek, "Remarks on the Algebra of Non First Normal Form Relations," Proc. of the ACM SIGACT-SIGMOD Conference on Principles of Database Systems, Los Angeles, Mar 29-31, 1982.

hd08 V. Linnemann, "Non-First-Normal-Form Relations and Recursive Queries: An SQL Based Approach," Proc. of the 3rd IEEE Intl. Conference on Data Engineering, Los Angeles, Feb 1987.

hd09 U. Arend, J. Wandmacher, "On the Generality of Logical Recording in Spatial Inference Tasks," Acta Psychologica 65, 193-210, 1987.

hd10 G. Rohr, "Mental Models - On the Relation Between Structure and its Symbolic Representation in Human- Computer Interaction," Zeitschrift fuer Psychologie 193, Heft 4, 1985.

hd11 G. Dueck, T. Scheuer, "Threshold Accepting - A General Purpose Optimizing Algorithm Appearing Superior to Simulated Annealing," to appear in SIAM Journal on Computing, 1989.

hd12 E. Keppel, et al., "Implementation of a Data Dictionary System to Maintain a Distributed Information System for Transplant Data," Lecture Notes in Medical Informatics 35, Springer-Verlag, 1988.

hd13 B. Alschwee, S. Grundmann, "System Design for a Computer Aided Juridical Expert System," In: A.A. Martino, F. Socci Natali (eds.): Automated Analysis of Legal Texts, Logic, Informatics, Law. Elsevier Science Publishers B.V. (North-Holland), 1986.

hd14 T. Bollinger, U. Hedtstuck, C.-R. Rollinger, R. Studer, "Text Understanding in LILOG," In: A. Kunz, U. Schmitz (eds.): Linguistic Approaches in Artificial Intel. Duisburg Papers on Research in Language and Culture, 6, Verlag Peter Lang, Frankfurt, 1988.

hd15 F. Guenthner, Hub. Lehmann, W. Schoenfeld," A Theory for the Representation of Knowledge," IBM J. of Res. & Develop. 30, No. 1, Jan 1986.

Houston, Texas (1966 - 1974)

ho01 George Paul, M. Wayne Wilson, "The VECTRAN Language," PASC Tech. Report G320-3334, Aug 1975.

ho02 L.B. Lesem, P.M. Hirsch, J.A. Jordan Jr., "Computer Synthesis of Holograms for 3-D Display," Comm. ACM 11, 661-674, 1968.

ho03 L.B. Lesem, P.M. Hirsch, J.A. Jordan Jr., "The Kinoform: A New Wavefront Reconstruction Device," IBM J. of Res. & Develop. 13, 150-155, 1969.

ho04 M.M. Adibi, P.M. Hirsch, J.A. Jordan Jr., "Solution Methods for Transient and Dynamic Stability," Proc. IEEE 69, 951-958, 1974.

ho05 T.A. Hughes, et al, "Application of Security Indices and Restraints to a Real Power System," Final Report to the Electric Power Research Institute, 1975.

Kuwait City, Kuwait (Established 1980)

ku01 S. Sami, "Towards High Quality Arabic Script Output," Panel on Non-Latin and Non-alphabetic Scripts, IFIP Congress 86: 10th World Computer Congress, Sep 1-5, 1986, Dublin, Ireland.

ku02 Y.A. El-Imam, "Synthesis of Arabic Speech using partial syllables," Proc. of the first Intl. Conf. on Arabization and Informatics, AED and IRSIT, Tunisia, Tunis, Mar 9-11, 1988.

ku03 O. Alameddine, "Arabseques and Koufi Art" Contribution in the book: "Die Pixel Revolution" published by IBM Deutschland GmbH, 1987.

ku04 A.R. Ragab, S.M. Metwalli, J. Rueda, "Preliminary Assessment of Image Processing Application to Large Deformation Measurement," Proc. of the Third Cairo University Conference on Mechanical Design & Production, Pergamon Press, Mar 1986.

ku05 S.M. Metwalli, A.H. Kamel, A.A. Saheb, "Surface Roughness Effect on Laser Speckle Spectral Density," SPIE's Technical Symposium Southeast on Optics and Optoelectrnoic Systems, Mar 31-Apr 4, 1986, Orlando, FL.

ku06 O. Alameddine, M.F. Quinn, K. Al-Jamal, "Experimental Measurements of the Downward Infra Red Sky Radiation in Kuwait," Applied Optics, 23, No. 3, Feb 1984.

ku07 K. Al-Jamal, O. Alameddine, H. Al-Shami and N. Shaabaan, "Passive Cooling Evaluation of Room Pond Systems," Solar and Wind Technology, 5, No. 1, 1988.

ku08 A.H. Kamel, S.E. Akashah, F.A. Leeri, and M. Fahim, "Particle Size Distribution in Oil-Water Dispersions Using Image Processing," in Computers and Chemical Engineering, 11, No. 4, 435-439, 1987.

ku09 M. Quinn, O. Alameddine et al, "A Deconvolution Technique for Determining the Intrinsic Fluorescence Decay Lifetimes of Crude Oils." Applied Spectroscopy, 42, No. 3, 1988.

Los Angeles, California (Established 1964)

la01 R.M. Buckley and W.B. Broadwell, "A Dynamic Model for Contouring NC Devices," in M. A. DeVries, ed., The Expanding World of Numerical Control, Village Press, New York, 156-161, 1972.

la02 E.H. Coughran, "A Preliminary Report on Shipboard Computing," Report G320-2615, IBM Los Angeles, 1968.

la03 W.J. Van De Lindt, "Movement of the Mohorvicic Discontinuity under Isostatic Conditions," Journal of Geophysical Research, No. 4 72, Feb 1967.

la04 B. Dimsdale, "A Natural Language Information Retrieval System," Proc. IEEE No. 54 12, 1966.

la05 K. Dahlgren, "Naive Semantics for Natural Language Understanding," Kluwer Academic Publ., Boston, 1988.

la06 M.M. Connors, C. Coray, C.J. Cuccaro, W.K. Green, D.W. Low, and H.M. Markowitz, "Distributed System Analysis: the Distribution System Simulator," Management Science 18, B425-453, 1972.

la07 B. Dimsdale, "On Multiconic Surfaces," IBM J. of Res. & Develop. 19, 523-529, 1975.

la08 R.M. Buckley and J.P. Mayfield, "Applications of a Numerical Geometry System in Engineering," Proc. IEEE 13th Annual Design Automation Conference, 25-33, 1976.

la09 P.S. Newman, "Towards an Integrated Development Environment," IBM Systems Journal 21, 1982.

la10 E.B. Fernandez, R.C. Summers, T. Lang, and C.D. Coleman, "Architectural Support for System Protection and Database Security," IEEE Trans. on Computer, C-27 8, 1978.

la11 E.B. Fernandez, R.C. Summers, and C. Wood, "Database Security," Addison-Wesley, Reading, 1977.

la12 R.C. Summers, "A Resource Sharing System for Personal Computers in a LAN: Concepts, Design and Experience," IEEE Trans. on Software Engineering, 895-904, 1987.

la13 A. Inselberg, "Multi-dimensional Graphics," Proc. of Eurographics '85, Nice, France, 315-325, 1985.

la14 A. Inselberg, "Intelligent Instrumentation & Process Control," Proc. of IEEE Conference on Artificial Intelligence Applications, 302-307, Dec 1985.

la15 A. Inselberg and B. Dimsdale, "Intelligent Manufacturing & Integrated Instrumentation," Proc. of ASME Conf. on Integrated Intelligent Manufacturing, Anaheim, CA, 341-357, Dec 1986.

la16 A. Inselberg, "Parallel Coordinates for Multi- dimensional Graphics," Proc. of Nat'l Comp. Assoc. Annual Conference, Philadelphia, 547-566, Mar 1987.

la17 A. Inselberg, "Cochlear Dynamics: The Evolution of Mathematical Model," SIAM Review, 20, 301-904, 1978.

la18 B. Dimsdale, "Convex Cubic Splines," IBM J. of Res. & Develop. 22, 168-178, 1976.

la19 S.A. Jurovics, "Solar Radiation Data, Natural Lighting and Building Energy Minimization," Proc. ASHRAE, 1979.

la20 D. Knight, M. Harry, J. Howard and J. Rivero, "Engineering Support Systems for Engineering Managers," IEEE Micro 5, 22-26, 1985.

la21 B. Arbab, "Object Identification from Parallel Light Stripes," Proc. 10th Inter. Joint Conf. on Artificial Intelligence, 1145-1148, 1987.

la22 G. Silverman, R. Tsai, and M. Lavin, "Locating Polyhedral Objects from Edge Point Data," Proc. of Intl. Joint Conf. on Artificial Intelligence (IJCAI), Milano, Italy, 1987.

la23 M.M. Parker, "Enterprise Information Analysis: Cost- Benefit and the Data Managed System," IBM Systems Journal 21, 1982.

la24 M.M. Parker and R.J. Benson with H.E. Trainor, "Information Economics: Linking Business Performance to Information Technology," Prentice Hall, Englewood Cliffs, 1988.

la25 P. Gongla and G. Sakamoto, "S*P*A*R*K: A Knowledge- based System for Identifying Competitive Uses of Information Technology," IBM Sys. Journal, IBM Systems Journal, Nov 1989.

la26 N.C. Shu, "Visual Programming," Van Nostrand Reinhold, New York, 1988.

la27 R.C. Summers, "An Overview of Computer Security," in Computer and Network Security, M.D. Abrams and H.J. Podel (eds), IEEE Computer Society Press, 9-25, 1987.

Madrid, Spain (Established 1972)

ma01 J. Brown, M. Alfonseca, "Solutions to Logic Problems in APL2," ACM SIGAPL, APL Quote Quad, 17-4, 356-361, May 1987.

ma02 M.L. Tavera, M. Alfonseca, J. Rojas, "The IBM Personal Computer APL System," ACM SIGAPL, APL Quote Quad, 14-4, 333-337, Jun 1984.

ma03 M. Udo, Y. Akimoto, S. Kaneko, T. Sanuki, M. Alfonseca," Japanese APL language system on IBM Multistation 5550," ACM SIGAPL, APL Quote Quad, 16-4, 326-334, Jul 1986.

ma04 L.F. Escudero and M. Rebollo, "A Mixed Integer Programming approach to Multispectral image classification," Pattern Recognition 9, 45-57, 1977.

ma05 A. Santisteban, L.R. Berriel and J. Bescos, "Image Restoration for a Defocused Optical System," Applied Optics 22, 2772-2780, 1983.

ma06 J. Jimenez, A. Santisteban, J.M. Carazo, J.L. Carrascosa, "Computer graphics display method for visualizing three-dimensional biological structures," Science, 232, 1113-1115, 1986.

ma07 L. Sopena, "Linguistica y Procesamiento de Textos en castellano," PC WORLD Espana, 8, 63-69, Feb 1986.

ma08 L. Sopena, "Natural Language Grammars for an Information System," ACM SIG Information Retrieval, 17, 4, 75-80, 1983.

ma09 J. Jimenez and M.A. Hernan, "Computer analysis of a high speed film of the plane turbulent mixing layer," J. Fluid Mech. 119, 323-345, 1982.

ma10 J. Jimenez, P. Moin, R. Moser and L. Keefe, "Ejection mechanisms in the sublayer of a turbulent channel," Phys. Fluids, 31, 1311-1313, 1988.

ma11 M. Aguilar, A. Garcia, P.J. Pascual, J. Presa and A. Santisteban, "Computer System for Scanning Tunneling Microscope Automation," Surface Science 181, 191-199, 1987.

ma12 L.E. Donate, L. Herranz, J.P. Secilla, J.M. Carazo, H. Fujisawa and J.L. Carrascosa, "Bacteriophage T3 connector: three-dimensional structure and comparison with other viral head-tail connecting regions," Journal of Molecular Biology, 201, 1988.

ma13 M. Kindelan, "Dynamic Modelling of Greenhouse Environment," Trans. ASAE 23, 1233-1239, 1980.

Mexico City, Mexico *(Established 1971)*

me01 J. Macias, M. Angulo, "A climatic-botanic information system of the state of Veracruz," Proc. XIII Conference of the Intl. Cartographic Assoc., Morelia, Mexico, Oct 1987.

me02 S. Castillo, E. Gutierrez, O. Alvarez, "Phenology of coastal dune vegetation," Proc. X Mexican Botany Congress, Guadalajara, Mexico, 1987.

me03 E. Gonzalez, S. Negrete, P. Noriega, "An evolving data base design," Proc. V Meeting of the Mexican AI Association, Merida, Mexico, 1988.

New York, New York (1964 - 1972)

ny01 M. Grigoriadis, W.F. Walker, "A Treatment of Transportation Problems by Primal Partition Programming," Management Science 14, 9, 565-599, 1968.

ny02 P.G. Comba, "A Language for Three-Dimensional Geometry," IBM System Journal 7, 3 & 4, 292-306, 1969.

ny03 Kurt Spielberg, "Algorithms for the Simple Plant Location Problem with Some Side Conditions," Operations Research 17, 1, 85-111, Jan-Feb 1969.

ny04 J. Greenstadt, "On the Reduction of Continuous Problems to Discrete Form," IBM J. of Res. & Develop. 3, 4, 355-363, 1959.

ny05 D.A. Quarles, K. Spielberg, "A Computer Model for Global Study of the General Circulation of the Atmosphere," IBM J. of Res. & Develop. 11, 3, 312-336, 1967.

Palo Alto, California (Established 1964)

pa01 H.P. Flatt, "Computer Modeling in Energy and the Environment," IBM J. of Res. & Develop. 25, 5, 571-580, 1981.

pa02 J. Gazdag, G. Radicati, P. Sguazzero, H.H. Wang, "Seismic Migration on the IBM 3090 Vector Facility," IBM J. of Res. & Develop. 30, 2, 173-183, 1986.

pa03 IBM Installed User Program 5796-PQW, "SORT/MERGE Performance Improvements."

pa04 A.A. Dubrulle, "A Version of EISPACK for the IBM 3090VF," Tech. Report G320-3510, IBM Palo Alto, Jun 1988.

pa05 Jeno Gazdag, H.H. Wang, "Concurrent Computing by Cascading Tasks in the Space-Time Domain," IBM Systems Journal, Nov 1989.

pa06 R.G. Scarborough, H.G. Kolsky, "Improved Optimization of FORTRAN Object Programs," IBM J. of Res. & Develop. 24, 6, 660-676, 1980.

pa07 IBM Installed User Program 5796-PKR, "FORTRAN H Extended Optimization Enhancement," 1978.

pa08 A.H. Karp, R.G. Babb II, "A Comparison of 12 Parallel Fortran Dialects," IEEE Software, 52-66, Sep 1988.

pa09 L.E. Lyon, "Striped I/O for FORTRAN," Tech. Report G320-3513, IBM Palo Alto, Jun 1988.

pa10 A.A. Dubrulle, R.G. Scarborough, H.G. Kolsky, "How to Write Good Vectorizable FORTRAN," Tech. Report G320-3478, IBM Palo Alto, Sep 1985.

pa11 R.G. Scarborough, H.G. Kolsky, "A Vectorizing FORTRAN Compiler," IBM J. of Res. & Develop. 30, 2, 163-171, 1986.

pa12 Hassitt, A., J.W. Lageschulte & L.E. Lyon, "Implementation of a High Level Language Machine," Comm. ACM, 16, 199-212, Apr 1973.

pa13 IBM Program 5799-ALK, IBM PRPQ MF2608 and IBM RPQ 500256, "APL/CMS and APL Assist," 1974.

pa14 A. Hassitt, L.E. Lyon, "An APL Emulator on System/370," IBM Systems Journal 15, 4, 358-378 1976.

pa15 H.J Myers, "Compiling Optimized Code from Decision Tables," IBM J. of Res. & Develop. 16, 489-503, Sep 1972.

pa16 IBM Installed User Programs: 5796-PJB, "APL Decision Table Compiler," 1977; 5796-PKD, "General Cross Assembler Generator," 1978; 5790-PPJ, "APL Performance Tools."

pa17 J. Littman, "The First Portable Computer," (Story of PASC's first PC, 1973), PC World, 1, 7, 294-300, Oct 1983.

pa18 J.V. Dave, "Contrast Attenuation Factors for Remote Sensing," IBM J. of Res. & Develop. 23, 214-223, 1979.

pa19 J.V. Dave, R. Bernstein, H.G. Kolsky, "Importance of Higher-Order Components to Multispectral Classification," IBM J. of Res. & Develop. 26, 712-723, 1982.

pa20 J.V. Dave, J. Gazdag, "Reduction of Random Noise from Multiband Data using Phase Relationships among their Fourier Coefficients," IBM J. of Res. & Develop. 28, 399-411, 1984.

pa21 IBM Program 5875-EEE, "Personal Computer Image Processing System (PCIPS)."

pa22 J.F.Asmus, R. Bernstein, J.V. Dave, H.J. Myers, "Mona Lisa," Perspectives in Computing, 7, 1, 11-22, 1987.

pa23 IBM Programs: PRPQ 5799-AYB, "Interactive Geo-Facilities Graphics Support"; 5668-941, "Graphics Program Generator Program Product"; "Graphics One-Line Diagram: Display, Edit, and Query Program (FDP)"; "Composed Document Viewing Utility" (Part of Graphic Data Query Facility, Rel. 3 (5668-905).

pa24 IBM Program 5796-PLN, "A Departmental Reporting System II (ADRS II)."

pa25 IBM Program 5798-DRT, "PROFS Note Maintenance Facility."

pa26 Paul J. Friedl, "The Use of Computer Conferencing by a University Faculty," IFIP TC6/WG6.5, Conf. on Message Handling Systems, 387-402, R. Speth, Apr 1987.

Paris, France (Established 1977)

pr01 R. Moreau, "Statistiques et Analyses Linguistiques," Presses Universitaires de France, Paris, 1966.

pr02 R. Moreau, "The Computer Comes of Age," MIT Press, Cambridge, (1984) (English Translation of "Ainsi Naquit L'Informatique," Dunod, Paris, 1981) Holland, Amsterdam 1968.

pr03 J.J. Duby, "Sophisticated Algebra on a Computer: Derivatives of Witt Vectors in Symbol Manipulation Languages and Techniques," North Holland, Amsterdam 1968.

pr04 F. Marcotorchino, P. Michaud, "Optimisation en Analyse Ordinale des Donnees," Masson, Paris, 1979.

pr05 F. Chatelin, "A Probabilistic Round Off Error Propagation Model," in Reliable Numerical Computing, Oxford University Press, 1988.

pr06 M. Abrams, A. Blusson, P. Nguyen, Y. Rabu, "Image Processing Applications for Geologic Mapping," IBM J. of Res. & Develop. 29, 2, 1985.

pr07 D. Ho, "A Soil Thermal Model for Remote Sensing," IEEE Trans. in Geoscience and Remote Sensing, 25, 2, 1987.

pr08 O. Raiman, "Order of Magnitude," Proc. of National Conference on Artificial Intelligence, AAAI-86, 1986.

pr09 P. Dague, P. Deves, O. Raiman, "Trouble Shooting: When Modelling is the Trouble," Proc. of National Conference on Artificial Intelligence, AAAI-87, 1987.

Philadelphia, Pennsylvania (1972 - 1974)

ph01 Franz Schober, Kurt Spielberg, "APL Econometric Planning Language (EPLAN)," Proc. of APL76 Conference, ACM, 1976.

ph02 Kurt Spielberg, "Inequalities in Branch and Bound Programming," in Optimisation Methods, (ed. Cottle, Karup), Engl. Univ. Press, London, 1974.

ph03 M. Guignard, K. Spielberg, "Logical Reduction Methods in Zero-one Programming (Minimal Preferred Variables)," J. of Operations Res., 29, Jan 1981.

Pisa, Italy (Established 1971)

pi01 A. Fahmy, L. Panattoni, E. Todini, "A Mathematical Model of the River Nile," Engineering Applications of Computational Hydraulics, 111-130, 1982.

pi02 C. Bianchi, G. Calzolari, P. Corsi, "A Program for Stochastic Simulation of Econometric Models," Econometrica, No.46, 235-236, 1978.

pi03 P. Franchi, C. Paoli, J. Gonzales, P. Mantey, A. Parolo, J. Simmons, "Design Issues and Architecture of HACIENDA, an Experimental Image Processing System," IBM J. of Res. & Develop. 27, No.2, 116-126, 1983.

pi04 C. Paoli, R. Porinelli, A. L'Abbate, M. G. Trivella, "A Computerized Autoradiographic Technique for Simultaneous Millimetric Mapping of Myocardial Blood Flow and Metabolism," IBM J. of Res. & Develop. 30, No.6, 627-634, 1986.

pi05 M. T. Pareschi, P. Armienti, G. Macedonio, "A numerical Model for the Simulation of Tephra Transport and Deposition: Applications to May 18, 1980 Mt.St.Helens Eruption," Journal of Geophysical Research, 93, No.B6, 6463-6476, 1987.

pi06 M. T. Pareschi, R. Bernstein, "Modeling and Image Processing for Visualization of Volcanic Mapping," IBM J. of Res. & Develop. 33, No.4, 406-416, 1989.

pi07 C. Paoli, R. Porinelli, C. Lettera, L. Masera, "Use of a Dictionary in Conjunction with Handwritten Texts Recognizer," IEEE Proc. of the 8th Intl. Conference on Pattern Recognition, IAPR, Paris, 699-701, 1986.

Brazil SC, Brasilia (1980 - 1986)

br01 M.A. Casanova, A.V. Moura, "An Introduction to Distributed Data Bases Management Systems," 1984.

Brazil SC, Rio De Janeiro (Established 1986)

ri01 M.A. Casanova, F.A.C. Giorno and A.L. Furtado, "Logic Programming and the Prolog Language," 1987.

ri02 R.C.B Martins, A.V. Moura, "Systematic Derivation of Correct Programs: The Denotational Approach," 1988.

Rome Italy (Established 1979)

ro01 F. Antonacci, P. Dell'Orco, A. Turtur, "A Territorial Database Management System," Proc. of Intl. Conference on Management of Data, ACM SIGMOD, Orlando FL, Jun 1982.

ro02 G. Volpi, G. Gambolati, L. Carbognin, P. Gatto, G. Mozzi, "Groundwater contour mapping in Venice by stochastic interpolation. 2 Results," Water Resources Research, 15, 3, 1979.

ro03 G. Fronza, P. Melli (eds.), "Mathematical models for planning and controlling air quality," Pergamon Press, 1982.

ro04 E. Runca, P. Melli, F. Sardei, "Finite difference and variational methods applied to numerical treatment of advection and diffusion of environmental pollutants from a line-source," Journal of Computational Physics, 59, 1, 1985.

ro05 L. Lenzini, G. Sommi, "Architecture and implementation of RPCNET," in Verma P.K. (ed.), Proc. of the Third Intl. Conference on Computer Communication, Toronto, 1976.

ro06 A. Fusi, G. Sommi, "Distributed Virtual Systems," Proc. 2nd Intl. Conference on Distributed Computing Systems, Paris, France, Apr, 1981.

ro07 S. Di Zenzo, R. Bernstein, S. De Gloria, and H.G. Kolsky, "Gaussian Maximum Likelihood and Contextual Classification Algorithms," IEEE Transactions on Geoscience and Remote Sensing. GE-25, 6, 1987.

ro08 S. Cavaliere, M. Fantini, A. Turtur "An Integrated System for Printing and Publishing Applications," in Computed-Generated Images, 436-447, Springer Verlag, 1986.

ro09 E. de Castro, G. Cristini, A. Martelli, C. Morandi, M. Vascotto, "Compensation of random eye motion in an electronic ophtalmoscope: preliminary results," IEEE Transaction on Medical Imaging, Mar 1987.

ro10 J. Gazdag, P. Sguazzero, "Migration of seismic data by phase shift plus interpolation," Geophysics, 49, 1984.

ro11 R. Benzi, H.R. Hansen, A. Sutera, "On stochastic perturbations of simple blocking models," Quarterly Journal of the Royal Meteorological Society, 110, 1984.

ro12 P. Carnevali, A. Selloni, A. Baratoff, E. Stoll, "Current distribution in the scanning vacuum tunnel microscope: free electron model," Journal of Physics C. Solid state Physics, 17, 1984.

ro13 F. Antonacci, M.T. Pazienza, M. Russo, P. Velardi, "A System for the Analysis and Generation of Sentences in Italian" in "Conceptual Graphs for Knowledge Systems," ed. J. Sowa and N. Foo, Addison Wesley, 1989.

ro14 F. Antonacci, M.T. Pazienza, M. Russo, P. Velardi, "Representation and Control Strategies for large Knowledge Domain: an Application to NLP," Journal of Applied Artificial Intelligence, 1989.

ro15 P. D'Orta, M. Ferretti, A. Martelli, S. Melecrinis, S. Scarci, G. Volpi, "A Speech Recognition System for the Italian Language," IEEE 1987 Intl. Conf. on Acoustic, Speech and Signal Processing, Dallas, 1987.

ro16 P. D'Orta, M. Ferretti, S. Scarci, "Fast Speaker Adaptation for Large-Dictionary Real-Time Speech Recognition," IEEE Workshop on Speech Recognition, Arden House, Harriman, NY, 1988.

Tokyo, Japan (Established 1970)

to01 T. Nishino and T. Fujisaki, "A Stochastic Parsing of Kanji Compound Words," Trans. of Information Processing Soc. of Japan, 29, 11, 1034-1042, Nov 1988.

to02 M. Udo, Y. Akimoto, S. Kaneko, and T. Sanuki, "Japanese APL Language System on IBM Multistation 5550," ACM/SIGAPL APL Quote Quad, 16, 4, 326-334, Jul 1986.

to03 J-K. Hong and J. Iisaka, "Coastal Environmental Change Analysis by Landsat MSS Data," J. Remote Sensing of Environment, 12, 2, 107-116, 1982.

to04 Y-L. Yang, A. Kondo and J-K. Hong "An Assessment of Environmental Changes by Landsat Data" in Ocean Space Utilization '85, Springer-Verlag, 1, 715-722, 1985.

to05 A. Koide, A. Doi, and K. Kajioka, "Polyhedral Approximation Approach to Molecular Orbital Graphics," Journal of Molecular Graphics, 4, 3, 149-156 & 160, 1986.

to06 M. Sakaki, H. Samukawa, N. Honjou, "Effective Utilization of the IBM 3090 Large Virtual Storage in the Numerically Intensive Computations of ab initio Molecular Orbitals," IBM Systems Journal 27, 4, 528-540, 1988.

to07 O. Wakita and S. Kadekawa "Computer for the Blind and the Deaf," Journal of IPSJ, 25, 5, 462-470, May 1984.

UKSC, Peterlee, England (1969 - 1979)

pe01 M.G. Notley, "The Peterlee IS/1 System," IBM Tech. Report UKSC-0018, IBM Scientific Centre, Mar 1972.

pe02 S.J.P. Todd, "The Peterlee Relational Test Vehicle - A System Overview," IBM Systems Journal, 15, no. 4, 285-308, 1976.

UKSC, Winchester, England (Established 1979)

uk01 M.L. Cocklin, "Digital Chest Radiology: Spatial and contrast resolution requirements," Proc. of the Intl. Symposium on Computer Assisted Radiology, Berlin, Springer Verlag, 215-219, 1987.

uk02 H.M. Cole, T.R. Harris, E.W. Stockley and H.V. Wheale, "Interactive Three-Dimensional Display of Reconstructed Neurons," Journal of Physiology, 400, 5P, 1988.

uk03 J.R. Woodwark, "Eliminating Redundant Primitives from Set-theoretic Solid Models by a Consideration of Constituents," IEEE Computer Graphics and Applications 8, 3, 38-47, May 1988.

uk04 G. Kaye, "The Design of the Database for the Survey of English Usage," Corpus Linguistics, Hard & Soft, Ed. M. Kyte, O. Ihalainen & M. Rissanen, Rodopi Amsterdam, 1988.

uk05 S.G.C. Lawrence, B.J. Williams and G. Kaye, "The Automatic Phonemic Transcription of English," Orthography and Phonology, Ed. P.A. Luelsdorff, J. Benjamins Amsterdam, 1987.

uk06 R.W. Phippen, G.R. Luckhurst and D.J. Adams, "Computer simulation studies of anisotropic systems XVII: the Gay-Berne model nematogen," Molecular Physics 61, 6, 1575-1580, 1987.

uk07 P. Reilly and R. Zambardino, "A computer-based system for the organization and processing of site-location and boundary data," Computers and Humanities 27, 177-191, 1987.

Venice, Italy (1969 - 1979)

ve01 G. Gambolati, "Use of over-relaxation techniques in the simulation of large groundwater basins by the finite element method," Intl. Journal of Numerical Methods in Engineering, 9, 1, John Wiley, 1975.

ve02 E. Runca, P. Zannetti and P. Melli, "A computer- oriented emissions inventory procedure for urban and industrial sources," J. of Air Poll. Contr. Assn., 6, 1978.

Washington, (Wheaton, Maryland) (1967 - 1969)

wa01 A. Hassitt, "Solution of the Stochastic Programming Model of Reservoir Regulation," Tech Report 320-3506, IBM Washington (Internal Publication), May 1968.

wa02 H.H. Wang, P. Halpern, P. Wrotenbery, "Primitive Equation Limited Area Fine Mesh Numerical Weather Prediction Model," Tech Report 320-3512, IBM Washington (Internal Publication), Oct 1968.

The Authors

Richard A. MacKinnon, *16 Page Farm Road, Sherborn, MA 01770-1034.*

Mr. MacKinnon is the third Manager of the Cambridge Scientific Center 1975-1992. He holds a B.A. degree in American History from Yale University (*magna cum laude*) and an M.B.A. in Finance from the Harvard Graduate School of Business Administration. He joined IBM in 1962, working first as a systems engineer then marketing representative in a banking territory in Boston. In 1966 Mr. MacKinnon installed the first on-line banking system on System/360. From 1967-1969 he worked in the banking development department in Data Processing Division headquarters, White Plains, New York. In 1970 he became Systems Support Manager in IBM's largest DPD region, New England and Upper New York state. There he was responsible for introduction and early support of System/370 and its virtual storage operating systems as well as associated hardware and system software products. In 1972 Mr. MacKinnon was responsible for the installation of the second System/370 shipped by IBM which was the first system installed using OS/360.

In 1975 Mr. MacKinnon joined the Cambridge Scientific Center as its Manager. Since 1976 he has taught the "Advanced Computing Systems" course at the MIT Sloan School of Management and currently holds the position there of Senior Lecturer. He is also a Visiting Lecturer at Harvard University's Aiken Laboratory. After closing the Cambridge Scientific Center on July 31, 1992 and retiring after thirty years of service with IBM, Mr. MacKinnon was appointed State Street Fellow upon the occasion of the two-hundredth anniversary of State Street Boston Corporation.

He is the author of four articles in the "IBM Systems Journal" on IBM multiprocessing, Virtual Machine operating systems, and the history of the Scientific Centers [ca10, ca11]. Outside IBM he is a member of the Foundation and the Board of Managers of the Massachusetts Eye & Ear Infirmary, a Harvard teaching hospital founded in 1824. He is also on the visiting committee of the radiation medicine department of the Massachusetts General Hospital as well as an elected member of Phi Beta Kappa.

Harwood G. Kolsky, *University of California, Santa Cruz, Computer Engineering, Santa Cruz, CA 95064*

Dr. Kolsky is a retired IBM Fellow, Palo Alto Scientific Center. He is a physicist who became a computer scientist when that field was new. His work has been on a variety of topics including programming languages, scientific applications, and digital image processing. He received his B.S. degree in Engineering Physics from the University of Kansas and his M.A. and Ph.D. in Physics from Harvard University.

After seven years at the Los Alamos National Laboratory, Dr. Kolsky joined IBM in 1957 in Poughkeepsie as a member of the product planning group for the STRETCH (IBM 7030) computer. In 1959 he became assistant manager of the IBM Federal Systems Division (FSD) office in Omaha, Nebraska. Following this he spent some time at FSD headquarters, before being named manager of the systems science department of the San Jose Research Laboratory in 1961. In 1962 he headed an advanced technology group in the Advanced Systems Development Division at Los Gatos, California. He joined the Palo Alto Scientific Center when it was formed in 1964 as manager of the atmospheric physics group. Later he headed projects in programming languages and digital image processing.

Dr. Kolsky was named an IBM Fellow in 1969. He served on the IBM Corporate Technical Committee at Armonk, New York, 1974-75. He was also head of the Board of Consultants for the IBM European Scientific Centers. In 1985 he came to the new UCSC Computer Engineering board of studies as a visiting professor. In 1986 he retired from IBM and began a new career as a full-time professor at UCSC. He is member of Sigma Xi, the American Physical Society, and a senior member of the IEEE.

Historical Appendix
Extracts of a 1973 Report
by
Harwood G. Kolsky

Historical Appendix

Extracts from a report entitled, " General Observations on the Role, Value and Problems of the Scientific Centers," by Harwood G. Kolsky, IBM Fellow, November 26, 1973.

III. The Effectiveness of the Scientific Centers

The DP [3] Scientific Centers serve two IBM needs in the real world of long range customer requirements: application research and scientific image building. In the opinion of independent auditors, the Scientific Centers serve these needs better than any other IBM organization. The reason for this is that the Scientific Centers are close enough to the real world to feel the pulse of customer requirements and to properly weigh technological trends that will influence DP sales. Yet they are sufficiently detached from the marketing organization so that the academic and scientific communities treat them as peers...

The overall quality of Scientific Center work is high. The productivity as measured by publications, application programs, external contacts and demonstrations is very high compared to other groups of similar size. The quality of Scientific Center work depends on having a critical mass of good people in several disciplines. The inter-disciplinary mix leads to maximum cross-industry activity. Small groups are forced to work with external people to get things done...

IV. Questions Facing the Scientific Centers

Whenever Scientific Center managers gather, certain problems are mentioned again and again as being important to the health of the Scientific Centers. Most of these problems are not the kind that can be "solved" once and for all by some management action. There are more like constraints which must be continually considered in all decisions...

A. Questions of Mission

The domestic Scientific Centers were established originally in 1964 and the WTC Centers started in 1968-69 to fill the needs of their respective countries in the gap between basic research and technology on one hand, and marketing development activities on the other...

Because of the three important aspects of the mission -- research, marketing, and image -- the Scientific Centers are always being pulled toward one of the other by management. The most successful projects have a proper balance between them. The concept that the Scientific Center manager should have an entrepreneur role in selecting things to do which will best serve the Center, science, and IBM's image and market, is not really understood very well, even by its own management...

The Scientific Centers give to IBM some of the advantages of a small company. Having small groups of sharp people, having long experience in their respective fields, running hard and fast on well-understood projects can result in incredible achievements. One could argue for having a few refuges

---

[3] DP = Dataprocessing Division. IBM's marketing division in the United States at the time.

within IBM for such activities on purely philosophical grounds, even with no immediate return on investment. The fact is that the return is better on a per-capita basis on almost any measurement which can be applied...

B. The Reporting Position Question

Because the mission of the Scientific Centers is balanced between research, market, and image, it is possible to have them report to any organization having one or the other of these missions in its charter. It can also report either at a high or low level in such an organization, with either a line or staff function.

The domestic Scientific Centers began reporting to Corporate Headquarters in a research, staff-like position, They were later shifted to DP, mainly for financial reasons. Their reporting point within DP has shifted several times over the years, indicating that there is uncertainty as to the "correct" position. The WTC Centers report either to DP or External Affairs (market or image) within their countries, although there is considerable variation in the exact conditions...

There is probably no single position which is "right", if by that one means no conflict will arise. The reporting point should be re-examined periodically to see if a change would help. In the final analysis, the Centers should report to that senior executive who is best willing to support and utilize them as SCIENTIFIC CENTERS, not as something else.

C. The Question of How to Allocate Resources

Financial problems always rate high on a Scientific Manager's list. Financial support is a concrete expression of the belief of IBM in a particular organization. Ii is interesting to ask how discretionary resources are assigned within IBM and match these against the Scientific Center's needs:

(1) The first is the formal way, which involves making marketing forecasts and building a business case before starting serious work on a project. This method works well on large projects which are extensions of IBM's existing product line. It does not work well on NEW opportunities, e.g., graphics, terminals, pocket calculators, or copiers. It also does not work well on small products with a limited market where conservative pricing policies can be deadly.

(2) Another approach is to find a project based on the inspiration of a charismatic leader. This works best on new opportunities previously overlooked by the company. The most spectacular advances come from such projects, but they can result in spectacular failures as well. The failures tend to generate more controls and more bureaucracy for future project leaders to face.

(3) One approach frequently used is that of estimating the return on investment while comparing two or three similar competing projects. Too often the resources go to the group which can put up the best P.R. story. The management problem is one of comparing unlike things. A "paper tiger" is always better than an existing entity, whether it is a hardware product, a software product, or the plans for a research project.

(4) Sometimes resources are applied on the basis of past performance. Intuitively one feels the best ways for IBM to assign resources would be to reward the people or groups who have a good record of accomplishment and punish those who have a poor record. All too often those who have squandered resources, missed milestones or produced shoddy products are rewarded by being given more resources to catch up, while those who are frugal and efficient are cut back. The Scientific Centers would certainly prosper if more decisions were made on the latter basis.

D. The Question of the Optimal Size for a Center

The really "optimal" size depends very much on the individual location, projects, personnel mix, etc. The closeness of coupling to neighboring universities and other scientific groups is particularly important...

E. Personnel Questions

There are two chronic personnel problems in Scientific Centers: (1) The hiring and retention of good staff members, and (2) the career path problem...

Because of the mission, Scientific Centers need staff members who (1) are good practicing scientists, (2) have a sense of the marketplace, and (3) have good judgment in dealing with other scientists, customers, and IBM management.

Having hired or developed these rare scientist- marketeer- statesmen, the problem inevitably arises as to what the career path is for such individuals. The same question occurs for scientists anywhere in IBM, but because of the small size of the Scientific Centers, it is particularly acute. This is clearly a continuing problem which must be faced on an individual basis...

F. The Questions of Project Selection and Monitoring

One of the most commonly asked questions is, "How are Scientific Center projects selected, monitored and killed?" A lot of time has been spent both in certain Scientific Centers and at headquarters explaining the procedures to be followed...

It has been observed in the Scientific Centers for some time that the really good projects come from the efforts of individual staff members who push ideas to fill a need which they can see because of their intimate knowledge of a field. Very few projects are ever handed down from above, or even suggested from Industry Development. Certainly none ever came through a formal flow chart of "criteria" and "decision points"...

G. The Question of Terminating Projects

The problem of terminating a project which is going nowhere is one of the more difficult management problems in the scientific environment. The arbitrary or capricious starting and stopping of projects will create lasting resentment and stifle individual initiative very quickly.

Too often the problem is approached from the point of view: "This project shall be killed to save resources, headcount and dollars." A much better way would be to ask, "What can we do which is more valuable with these particular people, with these particular skills, at this particular location?"

The best way to terminate a project is to start a better one. The people concerned will then switch over willingly rather than fight desperately to keep their old project alive. Stopping a well-run project simply because it has no short-term payoff is certainly contrary to the Scientific Center mission...

H. The Questions of Measurement

One of the oldest questions relating to scientists in industry is, "How can one measure the productivity of a scientist, a project, or a Center?" In the DP organization where everything tends to reduced to points, this is particularly important...

The secret certainty is people in the Scientific Centers. A project usually depends on one or two key individuals. This technical leader knows better than anyone else what the difficulties and the promises of the project are. As with our scientific customers, his judgment must be one of the key deciding factors in what the project is and how well it is doing.

To aid in the resource allocation problem and to avoid the problem of blind spots in projects, it is important to have technical and marketing reviews of the larger projects on a regular basis... If done properly, these reviews are more constructive than external "audits," which usually imply there is trouble with the project.

The method of these reviews is the same as the scientific method -- the submission of one's ideas to the scrutiny of his peers. This is the main way scientific progress is made and ideas propagated...

I. The Questions of Communication and Duplication

The bugaboo of "duplication of effort" ranks high in upper management's worries about projects in IBM. In the Scientific Centers it is hardly ever a real problem. Even projects having the same name are usually quite different in content or approach. As Manny Piore (Dr. Emmanuel Piore, former IBM Chief Scientist and head of Research) once said, "Duplication of projects is not necessarily bad. It is the lack of communication between similar projects which is unforgivable."

The best solution is to have staff members visit other Centers and give lectures on their work. Reports without personal contact tend not to be read or understood.

J. The Transfer Problem

The transfer of the results of scientific projects to other groups within IBM, or directly to customers, has been a continuing problem throughout the years. A project leader normally has to spend an appreciable part of his personal effort in seeking customers for his products. The percentage of success still remains questionable, since it is frequently hard to tell when a scientific "product" has been transferred. A concept used by another group can be very valuable although no paper transfer shows.

K. The Approval Procedure Problem

The difficulty of getting approvals...for new projects, joint studies, fellowships, special equipment, etc... separated from purely financial questions.

Scientific Centers tend to do unusual things, with the result that whatever they propose is likely to be a special case and tends to get stuck in the decision-making gears at all levels. Unlike some of the other problems on this list, this one does have some hope of solution as management and staff people become educated to the importance of Scientific Center projects.

★ ★ ★ ★

AFTERWORD:

A Final Summing Up
{July 1992}

Afterword: a Final Summing Up - July 1992

This report was originally written to honor the historical record of the IBM Scientific Centers on the occasion of their 25th anniversary in 1989. Its publication was delayed at that time because of the rapid succesion of events described herein. The authors now facing the reality that in 1992 all the U.S. Centers have closed, feel that it is our duty to complete and leave behind this historical record for all the Scientific Centers. In particular with their twenty-eight year history of accomplishment and contribution to IBM's objectives, the United States Scientific Centers and their members deserve that their legacy not be forgotten.

Why by 1992 did all the U.S. Centers Close?

First, it should be obvious from the previous pages that closures of Scientific Centers are not unique to 1992. Since the rapid growth of the mid-1960s, the number of United States Scientific Centers has been cut several times.

Part of this phenomenon was related to the growing expertise of customers in the application of computing to their endeavors. A pronounced emphasis arose in some minds within IBM and customer installations which indicated that IBM could best serve its broad marketplace through innovation in systems and computer science rather than through specific application expertise.

The result was that those Scientific Centers which were able to combine application knowledge with systems innovations flourished while the others were closed. Just as significant is the fact that, being part of IBM's United States marketing organizations, the Scientific Centers were subject to episodic and across-the-board resource cuts along with their marketing brethren. Belt-tightening was felt on a proportional scale in the Scientific Centers to a degree that was unappreciated by host organizational executives. In the authors' opinion, it is a tribute to the Centers' resourcefulness that they were able to endure so many years of such adversity.

By 1974 a so-called "steady state" was achieved with three remaining U.S. Centers. This level continued for 15 years.

This all began to change when the Los Angeles Center was closed in January 1990 as its parent organization, Technical Computing, applied significant resource cuts. It was concluded that for the overall long-term health of the Centers, two robust Centers in Palo Alto and Cambridge were preferable to three starving ones. Some -- but by no means all -- of the Los Angeles technical population transferred to Palo Alto. Thus in many ways, the Los Angeles scenario resembled the closings of earlier years. The intent in 1990 was to establish a new steady state in which Palo Alto and Cambridge were to remain in business and, in fact, they enjoyed robust and exciting technical agendas with enviable technology transfer achievements. This was true right up until the closing.

Perhaps over the prior twenty-six years the feeling had grown up that the Centers were "survivors" -- we knew how to be resourceful in difficult times and had acquired the secret for turning the pereception of keeping the "glass half full" into a reality while others saw it emptying. This combined with encouragement then tolerance from Corporate Headquarters gave a real basis for the conviction that the Centers' technical contributions were their best protection. Alas, this was not to prove so.

Instead, the remaining two Centers were terminated in 1992 because their parent organization, Technical Computing, could not be sustained and was abolished. In an IBM undergoing great adversity and change, the Centers came to be viewed as a $23-million expense rather than a seed investment which over many years had yielded rich dividends.

With the reversal of fortunes of Technical Computing, IBM executives first considered separating the Centers from the declining situation in which they were embedded, and sought a more tolerant and robust home in which the Centers could continue and hopefully flourish. Instead, in amazing haste, their candle was extinguished.

The IBM of 1992 -- referred to as "The New IBM" within the company -- has adopted a model that dictates that every activity beyond the Research Division must sustain a revenue flow and profit stream. When the Research Division declined to incorporate the Scientific Centers into its structure, Corporate Executives concluded there was no room for Scientific Centers in the IBM operation in the United States. [4] The authors have the satisfaction of debating this action at the highest levels of the IBM Corporation.

Fortunately such has not been the case outside the United States. The World Trade Centers continue to function and, while under considerable financial pressure -- mirroring their respective country's situations -- they continue and proudly remain a part of IBM outside the United States. Some countries such as Italy, Spain, and Israel have made considerable changes [5], but the Scientific Center mission, reality, and resource remains to a very large degree.

What to conclude at this juncture?

First, what was originally an organizational asset -- being part of an IBM domestic marketing organization -- became an ultimate liability and part of the downfall. When the Centers became a part of IBM's Applications Solutions Line of Business and Technical Computing operation, their horizon and time-line assumed a much fore-shortened perspective. The concept of investment for the future and the long-haul lost vogue, especially as business results turned unfavorable. Somewhere along the way the Scientific Centers lost their charter as a Corporate Resource and became treated as any other department within their host organization. This was coupled with the bewildering happenings in the mainframe and workstation marketplace -- the so-called "paradigm shift." Given the Centers' pioneering role with interactive then personal computing and their introduction of UNIX™ into IBM, one would think them an especially valuable resource at this juncture in IBM's affairs. But as the host endured rough seas and its ultimate demise, so went the Scientific Centers.

Lessons Learned

Under the category of *Lessons Learned* the authors feel one should never underestimate the extent to which an organization in the grips of adversity is willing to devour seed corn in the name of preventing starvation. Appeals to the future and references to "babies going out with bath water" fall on deaf ears under such circumstances. Appeals to previous contributions, sustained productivity, future importance, and enviable professional reputation also do not stay the adversity in such times as we faced in 1992. The "Endowment of Accomplishment" did not prove compelling to those making the decision.

---

[4] Private VNET communication from Mr. Jack Kuehler to H.G. Kolsky, March 22, 1992.

[5] Kuwait was never re-established after the 1991 Gulf War

Concluding Remarks

The Scientific Centers and their former members can be justifiably proud of their long record of service to IBM's customers on a wide range of scientific and technical projects. Because the Centers were a discretionary resource and allowed to work across the Corporation with Corporate oversight, they were able to undertake work which was occasionally counter to IBM's main strategies. Therefore, emphasis was placed upon openness and extensive communication of the technical activity with the objective of ***Center work becoming part of the strategy ultimately implemented by IBM.***

The authors look back and wonder at how prescient our founders were in 1964 -- not only in their choice of problems to be addressed but in their understanding of the challenging environment within which the problems would have to be solved. However, we are well aware -- as the material in the ***Historical Appendix*** from 1973 illustrates -- that many of yesterday's problems remained with the Centers throughout their long history. However, the productivity and creativity never stopped flowing. Results of the Centers' work have continued to be significant to the Corporation. The work represented a direct response to existing or emerging problems. Breaking new ground is just as important as addressing near-in problems.

Much has been written over the years on the difficulties and importance of technology transfer, and the Centers have a most positive story which we are here proud to tell. Equal importance, however, is attached to the need to understand and assess what is **not** going on within IBM and underway outside IBM. This spectrum of involvement -- within IBM and beyond -- is unique within the Company and helps account for the results achieved consistently over a relatively long time.

The Center's ability to work on real, significant problems, to affect IBM's technical agenda, and see its work flourish by practical use beyond the laboratory is what provided satisfaction to Center personnel. These satisfactions served as a positive counter-weight to the inevitable frustrations associated with long-term, demanding technical work which never had any guarantee of success or recognition. In retrospect we can see that the Centers ***did*** enjoy a success and professional recognition out of all proportion to their size and visibility. Such was the Centers' life within an industrial organization and our people both understood and thrived in it.

This report hopefully provides the reader with an appreciation for not only what the Centers have done, but by whom and where. From the authors' perspective, the United States Scientific Centers were entrepreneurs engaged in a unique and thrilling venture over twenty-eight years. Their legacy is a distinguished endowment of well-received and widely-used and highly-visible technologies for an incredibly modest investment. This, of course, is exactly what the founders hoped they would produce.

★ ★ ★ ★