The Centre for Advanced Studies: A model for applied research and development

by J. Slonim
M. A. Bauer
P.-Å. Larson
J. Schwarz
C. Butler
E. B. Buss
D. Sabbah

The Centre for Advanced Studies (CAS) is an applied research centre formed in 1990 within the IBM Toronto Software Solutions Laboratory. Its primary aim is to facilitate the transfer of research ideas into the various product groups of the laboratory. Although we are still learning how to make CAS operate more effectively, and it is too early to assess its long-term success, the model for CAS has proved to be workable. The primary partners, namely the IBM Toronto Software Solutions Laboratory, the IBM research community, universities in North America, and government agencies that support collaborative research, have found it a viable approach. As an overview, this essay provides some background to the formation of the centre, describes some of the challenges deemed important in defining the role of the centre, identifies a number of principles that are used to guide its formation and current operation, and reports on its progress. We conclude with a discussion of some lessons learned in the operation of the centre to date and identify future activities and directions for the centre.

The IBM Toronto Software Solutions Laboratory is a software development institution, a component of the IBM Software Solutions Division (SWSD).[1] SWSD has a worldwide mission for application-enabling software for all IBM computer platforms, and, increasingly, for other computer platforms in the open-systems marketplace. The application-enabling software breaks down into three logical businesses: database management, application development (including object-oriented software technology and the national-languages software architecture), and work group applications.

The role of the IBM Toronto Laboratory spans the development of strategic products and technologies for the database and application-development markets and for the national-languages architecture. It is one of the largest IBM software development sites and one of the largest private-sector research and development sites in Canada.

Evolving product lines and the multiple software and hardware architectures on which a product must operate introduce complexities in the development environment. Because the product and each of the architectures may be owned by different groups, numerous dependencies exist among groups. Thus, missions of the SWSD laboratories can be characterized as being diverse and interdependent. The development work is done at the sites having the best skills complement for a given assignment. Teams from different sites often cooperate on integrated solutions. For example, in database technology, the laboratory in Toronto works closely with the IBM Santa Teresa, Westlake, and Warwick laboratories; and in application-development technology, it works closely with the IBM Santa Teresa and Cary laboratories, and with the IBM German Software Development Laboratory in Boeblingen.

A significant shift has occurred within IBM from being a provider of hardware and software products to becoming a provider of solutions to customer problems. The laboratories within the SWSD are responsible for interacting with customers to identify problems, and for developing and maintaining products as well as integrating them into customer solutions. SWSD must also address the needs of customers within IBM. In particular, it provides enabling technology to the operating system groups that are responsible for the software environments of the different IBM computer platforms.

The Toronto laboratory has had very good experiences in cooperating with IBM research facilities. This cooperation led to some very successful product enhancements. For example, compiler optimization technology from the IBM Research facility in Yorktown Heights, New York, is incorporated in many Toronto compiler products in the form of TOBEY (Toronto Optimizing Back End with Yorktown). And Starburst database technology from IBM's Almaden laboratory in San Jose, California, is used in the Toronto DATABASE 2* (DB2*) products for Operating System/2* (OS/2*) and Advanced Interactive Executive* (AIX*) platforms. However, because the laboratory was limited to interaction with other IBM teams, additional influx of new ideas from outside IBM was needed.

Within the IBM organization, the acquisition of results from research in IBM is straightforward: there are few bureaucratic requirements, and groups are very cooperative. The research performed in IBM is first-rate and very attractive. However, even the capabilities of the IBM research laboratories are limited. The management at the IBM Toronto Laboratory recognized that additional resources and expertise were available from universities, various centres of excellence, customers, and government laboratories. The problem was how to effectively tap into this expertise and to become less isolated. Past relationships between IBM Canada and universities had been largely philanthropic, rather than promoting collaboration and the exchange of ideas. The challenge was how to establish such relationships. During the prior ten years IBM Canada had given $60 million to universities to support various programs, but it neglected to deal with individual researchers. Relationships with computer science and engineering researchers were effectively nonexistent. Paradoxically, the laboratory has acquired an increasing percentage of its employees from universities.

Several critical factors affected the operation and future directions of the IBM Toronto Laboratory.

Several critical factors affected the operation and future directions of the IBM Toronto Laboratory. First, the economic climate of recent years created a need for cooperation among different groups. Universities became much more willing to work with corporate partners. Because there are only a few sizable software research and development companies in Canada, the Toronto laboratory became an obvious choice. Second, IBM itself was in a transition from a proprietary computing environment to one of more openness and interoperability. Thus, the various hardware and software platforms prevalent in many university environments had become possible platforms for new and existing products. Third, the laboratory needed to become much more competitive by assessing the potential of new ideas and moving innovative ones from research to product as quickly as possible.

The previous and current directors of the IBM Toronto Laboratory saw great opportunities for the laboratory to grow. Their vision was to incorporate research results from within IBM and from external sources into successful products. They recognized further that a buffer was needed between development and research to facilitate the transfer of ideas between the groups. It had to be more than simply providing software professionals with research papers. At the same time, they did not want to create an exclusive research group. Hence, the idea of a research centre within the Toronto laboratory was born. The mission of the research centre was a simple one: to become a world-class applied research centre that facilitates the transfer of advanced research into strategic products. The applied research to be done must be defined by the customers, that is, the product groups within the Toronto laboratory.

Interactions between research and development communities is critically important in the transfer of research results.

The remainder of this essay describes the basic model of the organization established to carry out this mission: the Centre for Advanced Studies (CAS). Although many of the ideas and the organization of the centre are still maturing, the experience to date has been positive and should be shared. Many of the ideas are still incomplete, and it will be several years before the effectiveness of the centre can be thoroughly assessed.

The centre is not intended to replace any existing research facility within IBM; rather, it was predicated on taking advantage of these other groups and establishing a new group that works closely with the researchers at those facilities. Although the mandate of the centre is to perform applied research, we recognize the critical importance of basic research being done by other facilities and the fact that pure research must continue to be supported by both industry and government.

The next section identifies the challenges that the centre has faced and the opportunities that exist for it. Given these challenges and opportunities, the impediments and problems faced by the centre are then analyzed. An overview of the basic model for the centre is presented in a succeeding section, and areas in which we feel the centre has played a significant role are then discussed. Finally, we conclude with some lessons learned and directions for future activities.

Challenges and opportunities

The vision of an applied research centre that facilitates interaction between the research and development communities is predicated on bringing together individuals from both groups. Interactions between these groups has been identified as critically important in the transfer of research results within organizations, such as Hewlett-Packard Company and E. I. Du Pont de Nemours & Co., Inc.[2] Collaboration between the groups accommodates the transfer of research ideas to existing and future products. (See References 3, 4, and 5 for similar observations and additional comments.)

The notion of such a centre is not new. The Semiconductor Research Corporation,[6] for example, brings together companies involved in the semiconductor industry and university researchers. It has been successful in fostering research in this area, in training graduate students, and in transferring research results to its member corporations. However, the formation of such a centre focusing solely on problems in the software industry and devoted to the transfer of results in software research seemed to have few, if any precedents, and presented many challenges. The success of the centre ultimately depends upon addressing the challenges we now describe.

Enhance communication between research and development at all levels—The ultimate goal is an environment in which researchers and software professionals can interact directly. To achieve such an environment, we had to break the "perception barrier." The lack of contact and communication had resulted in distorted views: practitioners viewed researchers as "playing in the sand box" with no conception of reality, whereas researchers viewed practitioners as having closed minds, unable to grasp underlying concepts. Software professionals felt that researchers did not understand their problems and probably never had. Researchers, in turn, considered practical problems to be mundane and unchallenging.

Reduce time from idea to product—The time to transform a research idea into a successful product is often estimated to be 10 to 20 years. Shrinking this time is a fundamental challenge. Central to this challenge is selecting ideas ripe for transformation into products while avoiding ideas that are simply fads.

Applied research can help identify ideas with potential as new products or as parts of existing products. To be effective, the researchers must transform ideas into prototypes to validate their feasibility and uncover problems. Moreover, such prototypes must be developed in the context of actual systems, that is, integrated with parts of systems and other components. Stand-alone prototypes only address part of the feasibility issue; to be useful to professionals, the prototype must also demonstrate that the ideas can be integrated into an overall system architecture. Under such circumstances, professionals can then assess ideas and, when appropriate, make the move to transform prototypes and ideas into commercial products.

Exploit the professional's window of opportunity—The development and maintenance of a product is a cyclic phenomenon. A person's attachment to a product may be relatively short, only one or two years, before moving on to some other project. Software professionals are forced to react to the problems that arise during product development, giving them little time for innovation. There is usually a small window of opportunity during this time frame, consisting of the analysis and design stages, where new research ideas can be usefully incorporated. The challenge is to take advantage of this window of opportunity.

Accommodate the time frame of applied research—During any development project, problems arise that must be dealt with by the software professional. Unless there has been ongoing interaction with researchers, these release-specific problems are too short-term to involve them. These problems are often not anticipated during the analysis or design and so must be dealt with in a timely fashion. When more fundamental issues arise concerning a specific product (for example, increased functionality, performance, quality), they relate to releases one to five years in the future. In this case, there are more opportunities for researchers to become involved.

Applied research can help identify ideas with potential as new products or as parts of existing products.

Problems that require new ideas, concepts, and theories, such as migrating from a relational database technology to an object-oriented database or migrating from a sequential programming paradigm to a parallel paradigm, require even longer-term efforts. The challenge is to integrate research into long-term product development.

Transfer problems from development to research—Researchers produce solutions looking for problems, whereas professionals encounter problems looking for solutions. The difficulty is to match problems and solutions within a limited time frame. In research, there are well-known vehicles for exhibiting solutions, for example, papers and prototypes. In development, however, the questions are often not clearly formulated. The challenge is to make practitioners articulate these problems so that researchers can understand them. Research must be dovetailed with development at certain stages during the product life cycle. In the early stages, it is necessary to foster the professional's own awareness of such problems through aid from researchers. However, implementation and testing are best left solely in the domain of the professional.

Utilizing prototypes—The last 20 years have seen an exponential growth in the number of journals in computer science. Practitioners cannot find time to read and sift through all of this material. The challenge is how to spark interest among practitioners in new ideas and their possible commercialization. Prototypes illustrating concepts can be a useful vehicle to attract the interest of professionals. It suggests that encouraging re searchers to move their ideas from journals to prototypes might be a fruitful strategy.

Exploit external research sources—Although the internal research resources of IBM are excellent, only a limited number of problems can be explored. Thus, other sources must be considered. By making external research as easy to access as internal IBM research, we can explore additional problems. Concomitant with accessing external research sources is the challenge of facilitating the administrative mechanisms involved. This facilitation can be achieved, for example, by deemphasizing the role of formal agreements in the initial stages of interaction with outside researchers. Only when concrete results can be envisaged should legal aspects be handled. Essentially, administrative support services should become an asset in the overall process and not a liability.

The most important consideration was the value system of each group and the way in which each group measured success.

Increasing experts' awareness of products—The nature of computer science research, especially within the university environment, encourages researchers to focus on relatively narrowly defined problems within an area. This attention to detail often precludes individual researchers from developing an in-depth understanding of particular products associated with those concepts. For example, a database researcher may understand fundamental issues and concepts central to databases but have only limited knowledge of actual database products and their real-world use. The challenge is to find ways to turn these "experts" into individuals who understand and appreciate products related to their fields. In many cases, the researcher may have to forgo certain elegant solutions to accommodate product realities.

Accommodating differences in reward structures—Because of the different natures of the organizations, individuals in universities and industry have different value and reward systems (see the next section). A fundamental challenge is, first, to recognize these differences, and second, to provide a framework to accommodate the expected values and rewards.

Changing policies of funding agencies—In recent years, most university researchers have become aware of the increased emphasis on industrial collaboration by government funding agencies. However, if there are no mechanisms for encouraging such collaboration, or if no interest is shown in the results of the collaboration (including appropriate rewards), individuals and industrial partners will be reluctant to participate in collaborative efforts. The challenge is to make adjustments in the policies of interested groups (industry, university, government) to facilitate collaboration and enhance its benefits.

Recognizing and partitioning the risks—It should be the role of the industrial partner to take the financial risks and the researchers to take the intellectual risks. The challenge is to ensure that each group understands the risks it must accept.

Analysis

During 1990, members of the IBM Toronto Laboratory met several times with a small group of individuals from IBM Research and from universities. These meetings explored the possible nature of a centre for applied research. The framework of CAS grew out of the recognition that the most important needs of both researchers and software professionals had to be accommodated. The challenges and risks enumerated in the previous section were carefully examined in developing the paradigm for CAS.

Value systems. Based on our analysis, we recognized that the most important consideration was the value system of each group and the way in which each group measured success. Both systems engender a great deal of respect for individuals and their rights. However, there are also differences between the two cultures. Researchers place high value on:

- Scientific freedom, that is, the freedom of the researcher to pursue research problems of his or her own choosing

- Scientific recognition, that is, peer recognition for contributions to scientific knowledge

In contrast, practitioners value most highly:

- Market success
- Product excellence
- Career enhancements

In general, researchers aspire to breakthroughs within a discipline, whereas software professionals pursue products that are market leaders. The former are mainly concerned with conceptual issues; the latter are concerned with the details and challenges of a successful implementation.

Although these broad values act as a focal point for each group of individuals, others exist that are perhaps more specific. Using these values as a guiding principle, one can begin to identify the factors important to each group and those factors by which each group measures individual success. A number of such success factors are summarized in Tables 1 and 2.

Table 1 Success factors: Researchers

Recognized research impact
Peer recognition
International awards
Refereed publications
Patents
Research funding
Graduate students
Successful students
Individual monetary rewards
Citations

Table 2 Success factors: Software professionals

Product delivered on time and within budget
High-quality product
Individual productivity
Recognized innovations
Monetary rewards
Promotions
Respect from peer group
Respect as team member
Profitable products
Recognition as problem solver
Laboratory representative on task forces
Contribution to standards
Being on "fast track"

These tables contain only some of the success factors that motivate each of these groups. No attempt was made to be all-encompassing. Simply identifying and considering them in the development of the model for the centre was beneficial.

Impediments to cooperation. The initial analysis of the value systems and key success factors central to each group led to the identification of a number of impediments to cooperation. Table 3 summarizes the impediments and provides relative weights based upon their significance to each of the groups. A weight of 5 means most significant, and a weight of 1 means least significant. The weights are somewhat arbitrarily assigned based upon our experience. Clearly, the most significant were ones that the structure of the model had to address; eventually all will be addressed as we learn from experience and the model matures.

Some of the impediments in Table 3 are self-explanatory; we now discuss those that may not be obvious.

Lack of contact: There was no available (simple) mechanism for professionals and researchers to make contact with one another.

Buy-in of key personnel: Within any development group, often several individuals are perceived as leaders, wizards, or geniuses whom the other members of the group admire. These individuals could be managers, architects, designers, group leaders, or programmers; there is no specific pattern. Nevertheless, every successful group has at least one. These people sometimes perceive individuals from outside the group as threats to their own status within the group. Without their cooperation, the relationship of researchers with the group is difficult to develop and may not achieve as much success as possible. This problem is not only true of development groups; such individuals can also be found within research facilities as well. Conflicts between two such individuals can lead to long-term negative feelings.

Differences in problem scale: Some software professionals have the perception that researchers solve only "toy" problems because professionals are often forced to deal with large, complex products. Researchers may develop prototypes, but by their nature they are often designed and developed to avoid complexities so as to be able to

Table 3 Impediments to cooperation

Professional	Impediment	Researcher
5	Different value systems and focus	5
5	Lack of contact	5
5	Lack of buy-in of key personnel	4
4	Differences in problem scale	5
5	Accountability	3
3	Lack of trust	5
4	Lack of understanding of other's objectives	4
4	No perceived benefit to the individual	4
3	Lack of simple method for cooperation	5
3	Lack of senior management buy-in	4
3	Different computing environment	3
3	Differences in skill sets	3
1	Risk of discontinuity	5
3	Need for long-term relationships	3
3	No perceived immediate benefit to cooperation	2
2	Differences in jargon	2
1	Lingering student-teacher relationship	1

validate certain ideas. In contrast, a product can be thought of as a set of usable concepts.

Accountability: Both professionals and researchers take accountability seriously, though the mechanisms and impact are quite different. Professionals are accountable to those responsible for products, and their performance directly impacts the success of a product. Researchers are less directly accountable for their success to a single individual or group, but are, in a broad sense, accountable to colleagues and, in particular, to themselves.

Lack of trust: Experiences in the past for both groups left each with a number of misconceptions that needed to be corrected if cooperation was to be fostered. Practitioners felt that many researchers were or would be reluctant to become involved in the later stages of design and implementation. Researchers often felt left in the dark about how certain ideas, concepts, and techniques that they had developed had been used in products and often felt a lack of recognition for their efforts.

Lack of understanding of other's objectives: Again, there was a gap in recognition by both groups of those things most important to the individuals: for the researcher a good paper; for the professional a good product, among other objectives.

No perceived benefit to the individual: If any person involved in the relationship does not perceive any benefit, he or she will not make the effort to foster the relationship. Worse, if an individual sees goals being changed or values challenged—for example, a researcher not being permitted to publish research results—there is a disincentive to even begin a relationship.

Different environment: Typically, the hardware and software used by the respective groups differ. Perhaps more important, however, are the different organizational structures. Universities tend to be more open, operate more by consensus, and have few strict hierarchical reporting mechanisms. Within a corporate environment, the structure is more rigid.

Simple method for cooperation: Doing business with industry often entails lawyers, accountants, proposals, counter-proposals, etc. This process is generally time-consuming for both parties and sets the tone for the subsequent relationship as being adversarial rather than cooperative.

Differences in skill sets: Researchers are very involved in problem solutions, particularly in general settings; many problems often have no obvious solutions. The professional's skills are more disciplined, the solution is known, at least in broad terms, and the professional must translate that into a concrete implementation. A good skill for a practitioner, for example, is the production of simple, but elegant, code. A researcher may have greater skills at developing new, elegant algorithms.

Risk of discontinuity: The nature of the development environment is such that professionals move from group to group and project to project. The researcher, however, may not be able to follow the professional because the project may be outside the scope of interest or expertise of the researcher. As a result, the researcher is often at risk in spending the time and effort to develop individual contacts with practitioners who subsequently move on to other projects.

No perceived immediate benefit to cooperation: If the development team is facing a problem that requires an immediate solution and the researcher cannot provide a solution, the team sees little benefit, even though in the long term such cooperation might be extremely valuable.

Lingering student-teacher relationship: Many of the individuals working in the development groups graduated from universities. Software professionals are now being asked to treat university teachers as collaborators, whereas years earlier there had been a student-teacher relationship. There is some fear, albeit perhaps unfounded, that the "teacher" will tell them that they are "wrong."

Notice, in particular, the significant difference in weights of the two groups attached to several of these impediments in Table 3—in particular, the lack of trust, the risk of discontinuity, and the lack of a simple process for cooperation. All were considered very important in the context of the researcher and were much less important from the software professional's viewpoint. Again, these differences stem from the values of each group.

One of the most important aspects of this process was that little was formalized initially. Thus we were able to change our focus, views, directions, emphasis, etc., during discussions. Moreover, the analysis permitted us to view both sides rather than focusing on one or the other initially. As in any good relationship, both sides must be interested in communicating and must be committed to make it work and to "bend" on occasion.

The model

This section presents the CAS model by describing its fundamental principles, its underlying processes, and its initiatives to date. These fundamental principles are discussed in the following subsection. The CAS processes and initiatives to date are described in successive subsections.

CAS principles. The conclusion that an applied research centre must accommodate the values of both professionals and researchers led us to identify a number of fundamental principles that serve as the foundation for the centre.

Win-win design. An obvious principle of successful cooperation is a win-win design in which both parties gain something. This design requires understanding the main objectives of each party. Researchers need to understand that the main objective of a development organization is to find solutions to customer problems by obtaining or developing new technology that ultimately results in new products or improvements in existing products. As already noted, the main driving force for researchers is recognition and perceived research impact. This force has several implications. It is crucial that researchers are permitted to publish their results. It also means that university researchers are reluctant to work on issues that are proprietary or that, in their perception, will have little impact in the larger research community.

Mission-driven projects. All projects undertaken by CAS must be directly related to the missions of the laboratory and are in the short-to-midrange time frame. Preferably, they should originate directly from within a development group.

No permanent research staff. Having no permanent research staff in CAS is an important principle. Staff from development groups are assigned to CAS for specific projects for durations of six months to three years. The main reason is to make sure that the technology is rapidly and effectively transferred back into development. We

Figure 1 CAS research network

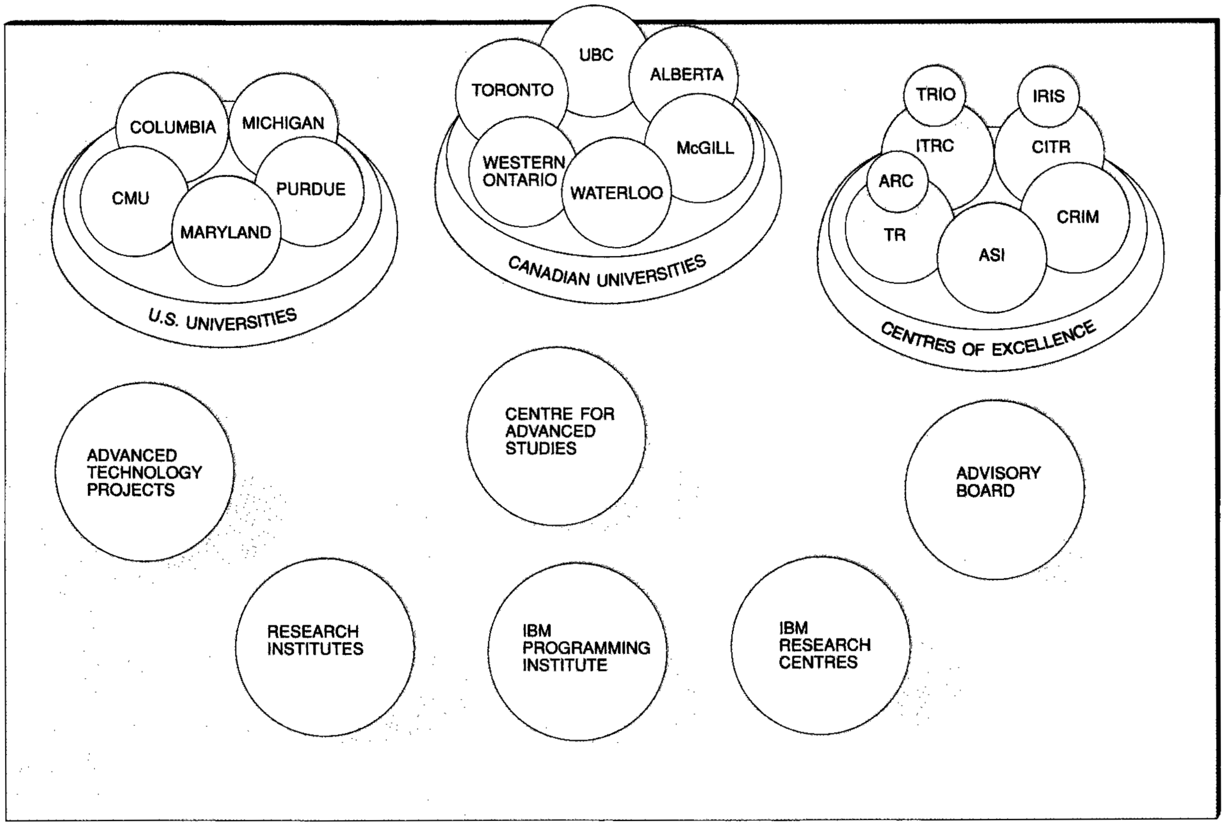

strongly believe that the best vehicle for transferring technology is people—people who understand the technology and who act as internal champions.

Projects funded by development groups. All projects must be partially funded by one or more development groups; typically they cover some of the salary expenses. This kind of investment from development groups is crucial to ensuring that the projects are relevant to the missions of the groups and that the groups "buy in" to the projects. If a project costs a group very little, there is little reason to consider it seriously or expect a return.

Foster simple, cooperative arrangements with experts. Project leaders are strongly encouraged to find and work with the best experts in the field, wherever they are located nationally or internationally—within the laboratory, IBM Research, universities, customers, other vendors, or government facilities. No one has a monopoly on good ideas. The mode of cooperation may differ, depending on the participants and the needs of the project. Flexibility in structuring the cooperation is important. One of the main functions of the centre is to facilitate building the team required for a project by locating experts and finding the best modes of cooperation. The centre actively maintains ties to various research groups and organizations throughout the world (see Figure 1). It follows that the scope, duration, and selection of participants need to be driven by the nature of the project, and not restricted by rigid rules.

Emphasize personal contacts and networking. The centre places heavy emphasis on personal contacts and joint work; personal contacts and working together builds trust. The benefits of personal contacts extend well beyond the duration of any single project. Simply being able to call some-

Figure 2 CAS external leverage strategy

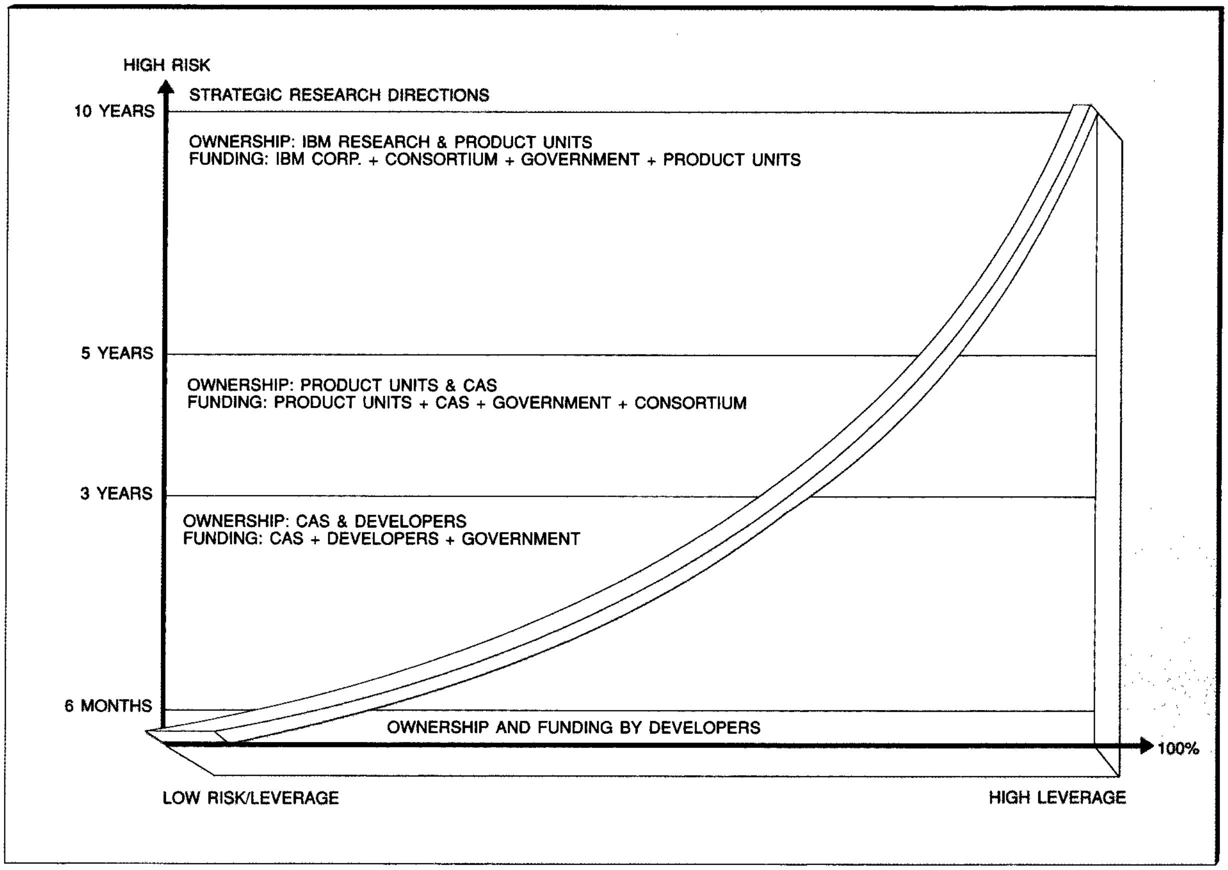

one to discuss a technical problem may sometimes be of great benefit for all parties. For professionals, in particular, it provides a wider source of expertise that they can tap into. Researchers, in turn, can gain insight into the internal aspects of large, complex software products. Moreover, transfer of technology is essentially a humanistic task;[2] without contact and trust, little can be accomplished.

Focus on prototypes. All projects so far are expected to bring the proposed technology to at least a prototype stage. The objective is to convincingly demonstrate the feasibility of the proposed technical solution; not only must it be conceptually sound, but it must also be implementable. This point is important; practitioners are not convinced by paper solutions. With prototypes it is easier to assess the implications of incorporating a technology within a product.

Leverage funding. Whenever possible, the funding directly provided by the laboratory is leveraged in a variety of ways to increase the scope of a project (see Figure 2). IBM has internal programs from which additional funding can be obtained. Governments in Canada and the United States also have various programs under which university participants can leverage funding obtained from industry. CAS focuses on low- to medium-term projects (six months to five years) with low-to-moderate risk associated with the research. Funding for higher-risk research is very important but is beyond the scope of CAS as an applied research centre.

CAS processes. With these principles in mind, CAS was built around a simple model of how to successfully develop and integrate innovative ideas into new or existing products. The model consists of a six-step process that encompasses the iden-

Figure 3 CAS project life-cycle process

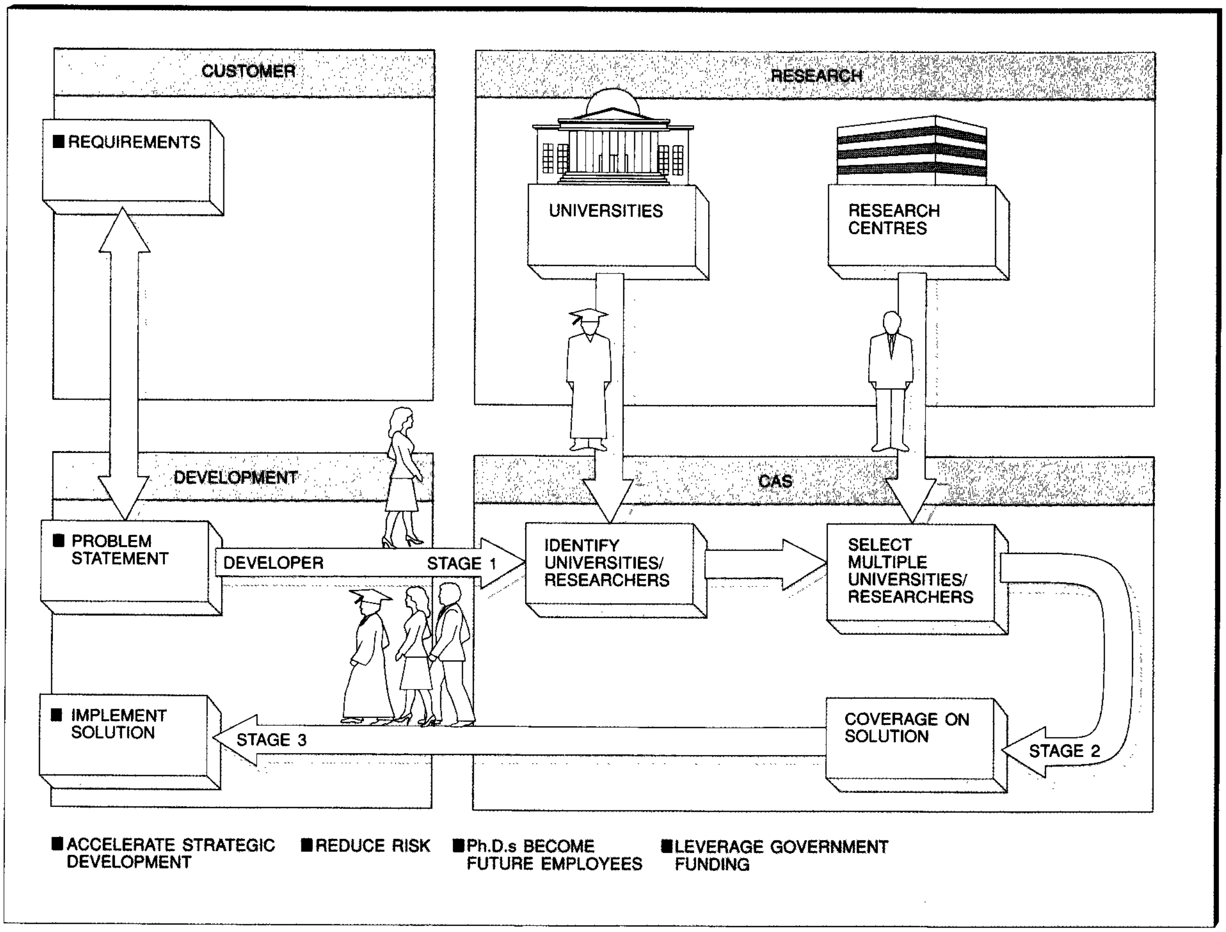

tification of problems, definition of a project, and creation of a project team. CAS projects follow a simple life cycle that consists of three main stages (see Figure 3). The steps within the basic CAS process are:

1. A new research project is initiated by one or more people in a development group.
2. The proposed project goes through an evaluation and selection phase.
3. Key people from interested development groups are assigned to CAS to work on the project for a specific period of time.
4. The project leader builds a research team to work on the project, tapping into the best sources of expertise in the area, wherever they can be found. The research done during this portion of the project is precompetitive; this is Stage 1 of the CAS project life cycle.
5. Once some viable technology has been identified and prototyped, the technology is evaluated by one or more development groups. At this point, the project has entered Stage 2.
6. At the end of the project, people assigned to the project are transferred back to their development groups to work on the new technology. This step represents Stage 3 of the project. Beyond this stage, the project may be terminated or may suggest the formation of new projects to pursue work on problems identified during the project or to investigate areas not adequately researched.

CAS initiatives. Based on the principles and basic CAS process, CAS has begun a number of initiatives aimed at fostering research to address problems central to the laboratory, transferring results to development and increasing communication among researchers and professionals. These initiatives are now described.

CAS projects (1990-1994). CAS has initiated 15 research projects to date. Some of these have already been completed and others are ongoing. Other projects are in the formative stages. More general information about some of these projects is provided in the next section.

Research fellowship program. The Ph.D. Fellowship Program is one of the cornerstones of the activities of the centre. Fellowships for graduate study have been key ingredients in other industry-university research centres.[6] CAS supplies university researchers with interesting problem topics related to their areas of expertise. Once agreement on a research topic to be explored has been reached, CAS provides scholarships for one or more Ph.D. students to work under a faculty member's supervision. The students pursue pieces of the selected problem as part of their thesis topic. At some point during their support, the students are expected to spend three or more months in CAS. This time permits them to network with other students, researchers, and professionals within the laboratory who are working in related areas. The result of the fellowship program is to speed up the research that a student must pursue. A major outcome of this program is that it is expected to increase the number of dissertations dealing with systems-related research topics and solutions to the identified problems.

Sabbaticals. For several years, the laboratory encouraged university researchers to spend sabbatical leaves at the laboratory, without much success. One of the main reasons for this lack of success was the difficulty that researchers had in interacting directly with development groups. As mentioned above, development groups are very much driven by short-term schedules and have little time for interacting with visitors. The existence of the centre and people in the centre working on research issues have made the laboratory much more attractive for sabbatical visitors. Shorter visits, for example, over a summer, are also encouraged.

Visiting scholars. Outside researchers working on joint projects are (strongly) encouraged to spend some time in the laboratory. The arrangements are flexible, ranging from one day a week to a few weeks or months at a time. We believe that personal contacts and actual working together are crucial to building trust and interaction beyond the lifetime of the project. In some cases, similar visiting arrangements have been made with leading researchers in some fields without direct involvement in a specific project. This method is regarded as an investment for the future, but it also has the immediate effect of reducing the gulf between practitioners and researchers.

University adjunct faculty. In the same spirit, laboratory staff are encouraged to become involved with (local) universities both as part of specific projects and as adjunct faculty members. In some cases, staff are teaching courses and supervising graduate students in their areas of expertise. Being involved in this way has the additional benefit of making students in the universities aware of IBM as a potential employer.

The Executive Advisory Board (EAB). The EAB is composed of the chairpersons of 12 computer science departments in Canadian universities that have Ph.D. programs, along with representatives from the IBM T. J. Watson and Almaden Research Centers. The EAB meets twice annually, and its mandate is to advise and review the mission of CAS. Recommendations about the goals and progress of CAS are reported to John Schwarz, Vice President, Application Development Solutions and Director of the Toronto laboratory. This board also acts as a forum in which both the concerns of Canadian universities and those of IBM Canada can be expressed. Concerns over the curriculum in computer science are often topics of discussion.

Research Partners Guide. This publication[7] lists the researchers with whom CAS has collaborated. Details are provided concerning the interests and background of each expert, as well as the current projects that are being undertaken in cooperation with CAS. The guide provides an excellent resource that permits software professionals to tap into a willing network of researchers without going through formal channels.

Technical reports. Every CAS project culminates in one or more IBM technical reports. The purpose is to provide useful documents concerning the project that can aid practitioners working in that area. Typically the generation of a technical report is also followed by a journal article or con-

Table 4 CAS projects and laboratory missions

Project	Distributed Services	Compilers, Languages	Databases	Programming Environments	Multimedia
Advanced software design technology				✓	
Distributed environment (CORDS)	✓	✓	✓	✓	
Color imaging processing					✓
Distributed knowledge worker	✓				
Extensible SQL/DS			✓		
Global on-line information					✓
Object-oriented environment				✓	
Parallel computing in C^{++}		✓		✓	
Program understanding			✓	✓	
Software processes				✓	
Ontario telepresence					✓
Multimedia services in high-speed networks	✓		✓		✓
Software reliability and testing				✓	
Multidatabases	✓		✓		✓
Managing distributed applications	✓			✓	

ference presentation that publicizes the contributions made. Because CAS members are drawn from different development groups within the Toronto laboratory, these individuals have increased both their expertise and their writing skills by the time that their CAS project is complete.

Centre for Advanced Studies Conference (CASCON). CASCON[8–12] is the annual conference sponsored by the centre. It provides a vehicle that brings researchers and practitioners together and allows them to communicate face to face. Besides providing information about the status of CAS-sponsored projects, it holds demonstrations of working prototypes. This event provides a sounding board for projects; constructive criticism from colleagues can be used to influence the future direction of projects. The conference permits researchers and professionals in different disciplines to come together and share ideas or to receive practical guidance from someone in a different area. Finally, CASCON permits the work of CAS to be visibly reviewed by both government granting agencies and IBM executives.

Impact of the centre

It is still too early to assess the full impact of CAS and to determine how effective it is in achieving a rapid transfer of technology. Nevertheless, it is still important to provide the reader with some assessment of how effective the model is.

To date, a total of 15 projects have been started; of these, nine are currently ongoing. A brief description of each of these projects and its relevance to the missions of the Toronto laboratory is presented in Table 4. A detailed summary of CAS projects can be found in the annual project summary technical reports.[13–15]

Based on our observations and discussions with others, we describe what we sense is the impact of CAS on the laboratory, on university research and education, and on government groups.

The IBM Toronto Laboratory. Impressions within the IBM Toronto Laboratory regarding CAS and its role are, as expected, both positive and negative. On the positive side, the increased communication between researchers and professionals is noted. On the negative side, CAS is still viewed as a "research" organization with all of the associated implications. In particular, this view creates expectations among professionals that the results of the work of the centre work should lead to the creation of new products that will generate new revenue for the laboratory. It is a mismatch between the perception that those people have and the current objectives of CAS that must be addressed.

We elaborate on the following areas in which we believe CAS has had an impact on the laboratory:

- One of the important objectives of CAS research projects was to reduce development risk by evaluating concepts and ideas via prototypes. Results of the projects may demonstrate that certain ideas are feasible and could be incorporated into products. As of today, the 15 CAS projects have yielded over 35 different prototypes. Of these prototypes, seven have embodied ideas or techniques of interest and have been transferred to development for further implementation. For example, concepts for novel distributed debuggers, graphical query languages, and parallel libraries in C++ are now being considered for commercialization.

  Research results can also demonstrate that our understanding of certain concepts is still immature or impractical and that their incorporation into products would be premature. Five prototypes proved to be in this category, including software for end-user calibration of colour images and multidatabases.

- Professionals have increased their personal contact with researchers. This increase has been achieved through involvement of professionals in the various CAS projects. The projects have involved around 120 faculty and around 160 graduate students.

- The CAS annual conference has been held for three consecutive years and has proven to be extremely important in bringing researchers, professionals, and others together. The past conference, CASCON'93, brought together about 900 participants in its four-day duration. Of these, approximately 40 percent were from universities and research centres, 33 percent from the Toronto laboratory, and the remaining attendees from other organizations, IBM Business Partners, customers, government groups, and other IBM sites.

- The visibility of CAS and the Toronto laboratory within the university community (individuals from over 50 Canadian and international universities participate in CAS projects) has helped recruitment for the laboratory. Graduate and undergraduate students have become much more aware of IBM as a preferred employer. This awareness is a significant turnaround from just a few years ago, when the Toronto laboratory was barely known among students. CAS initiatives have also increased the number of Ph.D. students and M.Sc. students, from both Canada and internationally, seeking employment with IBM.

- Through the interactions of researchers with professionals and the focus of CAS, more students get access to real-world problems and those important to the laboratory. As a result, these students have much more realistic views of industry needs, problems, and expectations.

- The use of university faculty within IBM's internal education facilities to teach courses and give seminars has also increased. It assists the laboratory in enhancing the skills of its professionals and in maintaining their technical vitality.

University research and education. The involvement of CAS with universities has also had an impact on university research and the education of graduate and undergraduate students in the following ways:

- Many of the CAS projects involve multiple universities; ten of the 15 CAS projects have involved more than three universities, and nine of these ten have involved at least one international university. Multiuniversity collaborative research makes it possible to attack larger and more complex problems. For example, the multidatabase project prototype contains code developed at four universities and also uses several commercial products. It now consists of over 260 KLOC (thousand lines of code) of university-generated code. Of course, this much code cannot be developed during the short time of a three-year project but represents the continuation of previous work. Continuation of work is precisely one of the goals of the collaborative research projects within CAS—to foster a research methodology that is not designed to start from "scratch" every time.
- Research funding from CAS and matching funds from governments has enabled university researchers in software and systems areas to build infrastructures and laboratories with people and state-of-the-art hardware and software to support research and education involving large, multisite projects.
- The increased openness to the IBM environ-

ment, via sabbaticals and student assignments at the laboratory, fosters a better understanding among students and faculty of development concerns. The increased exposure to product planning, to the internals of products, and to use of actual products as components in research results in better understanding of professional concerns. CAS has helped foster and encourage the growing trend toward experimental research in computer science and engineering by requiring that the feasibility of basic ideas be demonstrated and evaluated via prototypes.

- CAS has also had an indirect impact on university curricula. Through the efforts of CAS and faculty from several universities in Ontario, an M.Sc. program in computer science oriented toward software engineering has been established for laboratory employees. The program is unique since it is a joint effort by faculty from six different universities, and it sponsors courses given in intensive formats within the laboratory itself. CAS has also sponsored projects in testing, software processes, development environments, software architecture, program understanding, and other subjects.

Government funding agencies. The initiative by CAS to partially fund and support projects in software and systems areas with personnel and equipment has resulted in increased awareness of software research among the Natural Sciences and Engineering Research Council of Canada (NSERC) and the National Science Foundation (NSF) of the United States.

The CAS projects have helped to increase the relevance and potential of university research to the creation of wealth. Industry input to the research agenda ensures relevance and a higher potential for products and revenues.

CAS has been very successful in leveraging its investments in the projects. Within Canada, funds have been leveraged at a ratio of around one CAS dollar to five external dollars, and in the United States the ratio has been as large as one to ten.

By participating in government programs, software professionals were able to increase the government's understanding of the need for research in the software industry. Through flexible and responsive programs, in turn, government agencies made it more attractive for IBM to increase its funding of external software research.

Lessons, directions, and conclusions

There has been a great deal of learning during the initial four years that CAS has existed, and much work still lies ahead. Areas where we have gained new insights and adjusted expectations are briefly discussed below.

Stable funding is critical for the success of CAS since the model aims to establish trust based on long-term relationships between software professionals, researchers, and students earning a Ph.D. The CAS model is based on leveraging funds from external sources. Since current funding for the centre comes only from the laboratory, the financial pressure on the laboratory to decrease costs creates an unpredictable funding environment for the centre and associated researchers.

Our experience to date has shown that the success of CAS projects can be measured in different ways, and the criteria need to be established up front, on a project-by-project basis. The measurements of success are influenced by the different value systems of the groups. Increasing the understanding by both research and professional groups of the criteria of success of the other remains a challenge.

In CAS projects in which technology has been identified as useful, actually transferring it to development has proven to be as challenging as expected. The transfer process within the CAS model had not anticipated the need to synchronize technology knowledge transfer with the product planning cycle. Moreover, the success of the transfer was strongly correlated with the quality of the IBM champion of the project and the champion's respect within the laboratory.

Support or buy-in from the development groups in the laboratory is crucial to the long-term health of the centre. They are, after all, its "customers." Obtaining support requires sufficient communication between CAS and the development groups. It also requires the development groups to staff research projects within CAS with strong professionals, a critical need since the transfer of results from the project will depend on this individual. Experience to date suggests that staffing a project with two professionals enhances the chances of successful transfer.

Professionals should view working on a CAS research project as an opportunity and as a part of a normal career path. At this point, working on a CAS project is not integrated as a normal part of the career path of a software professional. It is seen by some as a career disruption or loss of opportunity for promotion and, therefore, not attractive.

Based on our experience, we can also draw some conclusions about CAS and about aspects of collaborative projects involving researchers and professionals. A common research environment is necessary. It helps foster common terminology and contexts when discussing problems, solutions and examples. The problems must be challenging so that the researchers see some benefit in the overall research effort. Financial support for Ph.D. students as a vehicle for supporting collaborative projects has been effective. Finally, intellectual property issues have not been as great a problem as expected. CAS projects have published over 200 articles and have also yielded more than ten patent filings. When professionals and researchers cooperate and try to accommodate the needs of each other, difficulties can be overcome.

In summary, our analysis reveals that there are significant differences in the objectives of practitioners and researchers. Recognition and acceptance of this fact are the keys to successful collaboration. Each group has to understand the objectives of the other. It follows that collaborative programs must be structured so that the objectives of each party are at least partially met.

Different objectives and lack of communication and interaction creates distrust, misunderstanding, and lack of appreciation. We strongly believe that increased communication, interaction, and openness are crucial to building trust and understanding. Furthermore, one of the best ways to achieve this goal is for people to work together, in the same location, on the same problem, over a period of time. The importance of personal relationships should not be underestimated. However, it is necessary to be realistic. Building trust and understanding takes time. It has to be an ongoing process.

The papers in this issue of the *IBM Systems Journal* describe some of the CAS projects.

Acknowledgments

The authors would like to acknowledge the support and input of the many individuals who contributed to the establishment of the Centre for Advanced Studies and its ongoing operations. From IBM, the authors would like to thank Larry Achtemichuk, Nick Cooper, John Botsford, John Henshaw, Pat Finnigan, Ron Holt, Arthur Ryman, Earl Wheeler, Steve Mills, Abe Peled, Shaula Yemini, and Pat Selinger for their individual contributions. From universities and research centres, the authors acknowledge the help and support of Maria Klawe, Hausi Müller, Ric Holt, Al Borodin, Ian Munro, Mike Leveson, Renato De Mori, Eshrat Arjomandi, David Taylor, Kellogg Booth, Morven Gentleman, Toby Teorey, and Yechiam Yemini. A special thanks goes to Ann Gawman for contributing to the preparation of the papers in this issue within a very brief period of time. There are so many others who contributed to the centre that it is not possible to list all their names here; we hope they understand.

Cited references and note

1. The IBM Software Solutions Division includes development sites in San Jose, California; Cary, North Carolina; Westlake, Texas; Boeblingen, Germany; Vienna, Austria; Warwick, England; and Dublin, Ireland, and software manufacturing sites in Boulder, Colorado; Charlotte, North Carolina; and Copenhagen, Denmark.
2. G. Bylinsky, "Turning R&D into Real Products," *Fortune*, 72–74 (July 2, 1990).
3. Computer Science and National Research Council Technology Board, *Keeping the U.S. Computer Industry Competitive*, National Academy Press (1989).
4. R. Tummalo and F. Cabrera, *University Relations in a Changing Environment*, IBM Academy of Technology, IBM Corporation, Armonk, NY (1994).
5. Government University–Industry Research Roundtable: Academic Industry Program, *New Alliances and Partnerships in American Science and Engineering*, National Academy Press (1986).
6. R. Cavin, L. Sumnet, and R. Burger, "The Semiconductor Research Corporation: Cooperative Research," *Proceedings of the IEEE* 77, No. 9, 1327–1344 (September 1989).
7. L. Baron, *Centre for Advanced Studies Research Partners Guide*, IBM Canada Ltd. Laboratory, Centre for Advanced Studies, Toronto (1992).
8. *Proceedings of the 1992 CAS Conference (CASCON'92), Volume II*, M. Bauer, J. Botsford, P. Larson, and J. Slonim, Editors, IBM Canada Ltd. Laboratory, Centre for Advanced Studies, Toronto (1992).
9. *Proceedings of the 1992 CAS Conference (CASCON'92), Volume I*, J. Botsford, A. Ryman, J. Slonim, and D. Tay-

lor, Editors, IBM Canada Ltd. Laboratory, Centre for Advanced Studies, Toronto (1992).
10. *Proceedings of the 1993 CAS Conference (CASCON'93), Volume I*, A. Gawman, M. Gentleman, E. Kidd, P. Larson, and J. Slonim, Editors, IBM Canada Ltd. Laboratory, Centre for Advanced Studies, Toronto, and National Research Council Canada (1993).
11. *Proceedings of the 1993 CAS Conference (CASCON'93), Volume II*, A. Gawman, M. Gentleman, E. Kidd, P. Larson, and J. Slonim, Editors, IBM Canada Ltd. Laboratory, Centre for Advanced Studies, Toronto, and National Research Council Canada (1993).
12. *Proceedings of the 1991 CAS Conference (CASCON'91)*, A. Gawman, J. Pachl, J. Slonim, and A. Stilman, Editors, IBM Canada Ltd. Laboratory, Centre for Advanced Studies, Toronto (1991).
13. J. Botsford, *Centre for Advanced Studies 1993 Project Summaries*, Technical Report TR 74.136, IBM Centre for Advanced Studies, Toronto (1994).
14. J. Henshaw and A. Ryman, *Centre for Advanced Studies 1992 Project Summaries*, Technical Report TR 74.088, IBM Centre for Advanced Studies, Toronto (1992).
15. A. Ryman and R. Helms, *Centre for Advanced Studies 1991 Project Summaries*, Technical Report TR 74.084, IBM Centre for Advanced Studies, Toronto (1991).

Accepted for publication April 15, 1994.

Jacob Slonim *IBM Software Solutions Division, Centre for Advanced Studies, Toronto Laboratory, IBM Canada Ltd., 844 Don Mills Road, North York, Ontario M3C 1V7, Canada (electronic mail: jslonim@vnet.ibm.com).* Dr. Slonim is head of research for the Centre for Advanced Studies of the IBM Toronto Software Solutions Laboratory. He received his Ph.D. from Kansas State University in 1979. He is an adjunct professor at the University of Western Ontario and the University of Waterloo. He has served on several international standards committees and has helped to organize numerous international conferences. In 1993 he became a member of the IBM Academy of Technology. Dr. Slonim was project leader for the CORDS project. His research interests include databases, distributed systems, and software engineering.

Michael A. Bauer *Department of Computer Science, University of Western Ontario, London, Ontario N6A 5B7, Canada (electronic mail: bauer@csd.uwo.ca).* Dr. Bauer is chairman of the Computer Science Department at the University of Western Ontario. He holds a Ph.D. in computer science from the University of Toronto. He was a principal investigator in the CORDS project within CAS. His research interests include distributed computing, distributed directories, and software engineering.

Per-Åke Larson *Department of Computer Science, University of Waterloo, Waterloo, Ontario N2L 3G1, Canada (electronic mail: palarson@maytag.uwaterloo.ca).* Dr. Larson is a professor in the Department of Computer Science at the University of Waterloo. He served as chairman of the department during 1989–1992. He was a principal investigator in the CORDS project within CAS. His research interests include multidatabase systems, query optimization and processing, and parallel databases.

John Schwarz *IBM Software Solutions Division, Toronto Laboratory, IBM Canada Ltd., 895 Don Mills Road, North York, Ontario M3C 1W3, Canada (electronic mail: needs@torolab2.vnet.ibm.com).* Mr. Schwarz is Vice President, Application Development Solutions and Director of the Toronto laboratory. He joined IBM in 1974 in Manufacturing Information Systems and joined the Toronto laboratory in 1984. In 1992 he was named laboratory director, leading a 1300-person team with worldwide missions for application development and data solutions. He assumed the additional responsibility for Application Development Solutions in the Software Solutions Division in December of 1993, which encompasses a team of approximately 1200 professionals across four development locations worldwide who are engaged in application development solutions for IBM.

Cindy Butler *IBM Software Solutions Division, Centre for Advanced Studies, Toronto Laboratory, IBM Canada Ltd., 844 Don Mills Road, North York, Ontario M3C 1V7, Canada (electronic mail: butler@torolab2.vnet.ibm.com).* Ms. Butler is the University and Research Project Coordinator for the Centre for Advanced Studies in the IBM Toronto Software Solutions Laboratory. She has 20 years of administrative experience with IBM in several areas. She has spent the past four years working with universities to foster and develop relationships with university administration, faculty, and graduate students. Her main responsibility is the Ph.D. scholarship program, which includes coordinating and monitoring visits and the progress of the graduate students. In 1993 she received special recognition for her contribution to this program.

Erich B. Buss *IBM Software Solutions Division, Software Engineering Process Group, Toronto Laboratory, IBM Canada Ltd., 895 Don Mills Road, North York, Ontario M3C 1W3, Canada (electronic mail: buss@vnet.ibm.com).* Mr. Buss is an Advisory Software Engineering Process Analyst in the Software Engineering Process Group of the IBM Toronto Software Solutions Laboratory. He graduated with an M.Sc. in computer science from the University of Western Ontario in 1976. He joined IBM in the SQL/DS data group in 1988 and subsequently moved to the Centre for Advanced Studies in 1990. In CAS he was the IBM Principal Investigator for the Program Understanding Project, where he worked on the practical application of reverse engineering technology to real development problems. His current interests are in program analysis, defect filtering, and object-oriented development.

Daniel Sabbah *IBM Software Solutions Division, Route 100, Somers, New York 10589.* Dr. Sabbah is the Director of Application Development Architecture for IBM's Software Solutions Division. Prior to that, he was Director of Software Technology in IBM Research. He began his career in IBM in 1974 in telecommunications software (VTAM) in Kingston, New York. He received his Ph.D. from the University of Rochester in 1981, specializing in artificial intelligence and computer vision. He then returned to the IBM Research Division and was responsible for restarting and increasing the growth of the artificial intelligence effort until 1987. He subsequently took over the programming languages effort and then software technology. He had overall responsibility for work in computer languages, compilers, programming environments, and software engineering. Dr. Sabbah has direct experience in product development and in software research.

Reprint Order No. G321-5547.