IBM's directions in technical computing

by P. L. Prairie
A. H. Weis

Technical computing comprises hardware systems, software, tools, communications networks, and applications to significantly increase productivity and competitiveness. A Technical Computing Structure (TCS) is described which provides a framework and IBM's direction to integrate these elements, including both IBM offerings and industry standards to support the technical end user. Also discussed are some of the special studies undertaken to improve development programs for this environment.

Computing is an important enabling technology for both competitiveness and productivity and has a pervasive impact on both individual enterprises and entire industries at a national level. In many cases, the computing environment manifests itself as two distinct types: One encompasses the data processing facility of the enterprise, which frequently provides the underpinnings on which the entire business is dependent. The other encompasses the research and development community within the enterprise, typically used for technical computing and engineering/scientific purposes. Driven by enormous needs for computing power, technical computing frequently pushes the state of existing computer technology in terms of both capacity and function.

This continuing need for more performance and function, along with the rapid advent of advanced computer technology, have fostered explosive growth in the application of that technology to research and development. Whether this research or development is done by technical users in an industrial/commercial environment, scientists in an industrial environment, or research scientists in an academic environment, it leads all of these groups to share similar computing needs. These are summarized as follows.

1. The need to select the best hardware and software for each aspect of the workday to minimize the time it takes to solve a problem or perform a task. The selection is usually dependent on function, performance, and/or price characteristics. The choice of equipment frequently consists of both workstations, which are dedicated to individual users and provide the user interface and local computing resource, and large systems, which are shared among the community of users, providing global compute and data facilities. Moreover, given the explosive growth of computing power and function within the industry and the wide range of participating vendors, this choice of systems is frequently divided among multiple vendors to a much greater degree than is usual in the data processing portion of the installation.
2. The ability to interconnect hardware and software products into a network of peers to take advantage of the strengths of the individual com-

© **Copyright** 1988 by International Business Machines Corporation. Copying in printed form for private use is permitted without payment of royalty provided that (1) each reproduction is done without alteration and (2) the *Journal* reference and IBM copyright notice are included on the first page. The title and abstract, but no other portions, of this paper may be copied or distributed royalty free without further permission by computer-based and other information-service systems. Permission to *republish* any other portion of this paper must be obtained from the Editor.

ponents. Typically, the network establishes client/server relationships where the workstations act to present information for review, update, and analysis of computational results, and the large shared systems provide computational support, database, file, information, and network services to other systems and networks.

3. The incorporation of vendor-supplied application solutions, as well as internally developed applications.

4. A consistent and friendly end-user interface available on technical workstations and personal computer systems. In many cases, end-user productivity for technical professionals is considerably enhanced by the ability to visualize both the results and the methodology used in the computation interactively. This phenomenon, known in the industry as the visualization of scientific computing, or *visic*, has accelerated the advance of workstations with graphical user interfaces between the technical end user and the networked computing environment.

   Often, these workstations run a variant of the UNIX® operating system. The UNIX system has enjoyed tremendous success in the scientific and academic communities, in contrast to its recent and still emerging presence in data processing environments. This success is due in part to the historical evolution of the UNIX system as a tool for technical professions and in part to its success as a vehicle for application development. But its success is due particularly to its ability to support the heterogeneous systems environment which proliferates in the technical computing community of most enterprises.

5. The ability to utilize numerous industry standards and evolving standards to improve the technical professional's productivity and effectiveness.

6. The ability to rapidly move proven methods and applications to production science or engineering environments across an enterprise, discipline, or industry.

This paper discusses the characteristics of IBM's approach to technical computing, including goals and some special studies and business relationships established to gain insight and accelerate our offerings in this important area. IBM has added substantial structure and organization to its technical computing offerings and is referring to this environment as the *Technical Computing Structure.*

Goals

IBM's goals in technical computing are to provide a framework for technical professionals to perform their jobs, manage their workday, and work within an integrated enterprise. This framework encompasses application development, computational analysis, and services such as office and technical publishing, and includes suites of vendor applications. The prime goal of the Technical Computing Structure is to establish a computing environment that minimizes the time to obtain the solution to the technical computing user's problem.

TCS has a synergistic relationship with IBM's System Application Architecture (SAA).

Technical Computing Structure

The Technical Computing Structure (TCS) provides an environment where industry standards are incorporated into IBM products, where current products are extended with a focus on key technical computing requirements, and where effective bridges are provided between IBM and non-IBM product offerings.

TCS has a synergistic relationship with IBM's Systems Application Architecture (SAA) and interoperability in a cornerstone of that relationship. TCS users will be able to take advantage of a growing number of IBM's SAA offerings such as database, network management, and integrated solutions. Such offerings already exist in several areas: For example, the National Science Foundation's Transmission Control Protocol/Internet Protocol (TCP/IP)[1] network is managed by IBM's NetView™. Enterprises that have large SAA networks can route TCP/IP traffic over their Systems Network Architecture backbones, and gateways are available between the TCP/IP mail servers (SMTP)[2] and PROFS™ (IBM's Professional Office System).

TCS is founded on end-user task analysis. On the basis of that task analysis, the systems, applications, and product strategies are formed, and selected associations are established with customers, universities, and software vendors. Early introduction programs are identified to bring quality products that are timely to the marketplace in a controlled environment.

The components of TCS describe an evolving integrated and seamless computing environment for the technical professional.

We have identified several representative application areas for TCS. These areas include, but are not limited to, Computer Integrated Manufacturing (CIM), Design Automation (DA), Computer-Aided Software Engineering (CASE), chemistry, and the process industry. In each of these representative application areas, scenarios have been developed to assess the tasks performed by technical professionals.

For example, in CIM, one scenario is the process of producing specifications for a new part. At all times, it must be possible to obtain a graphic display of the relationships between the data describing a part and the many versions of data that are produced during an interactive design process. In addition, structural information, such as the defined hierarchy of data that comprises an assembly, must be accessible and able to be visualized. This scenario depicts the complexities and data relationships associated with product and component design through release to manufacturing. It affects computational capability, end-user interface, data management, network management, and applications.

Description of TCS

The components of TCS describe an evolving integrated and seamless computing environment for the technical professional. The result is a network of workstations and servers, as shown in Figure 1. Figure 1 represents a conceptual bridge linking workstations and servers. TCS is represented as layers on this bridge. The user interface is often a UNIX workstation, with support for industry standards and non-IBM interconnections. The servers will provide computation, database, file, and information and program distribution functions.

IBM's direction, to achieve our objective of minimizing the time to solve technical computing problems, is to provide a rich set of offerings in five major areas: performance, industry standards, networking, systems and tools, and applications.

Performance. A major emphasis is being placed on all aspects of performance. It includes, but is not limited to, faster workstations, supercomputers, and other server systems. In addition, there will be continued enhancements to the related systems and software to capitalize on higher levels of vectorization, parallelism, larger memories, and higher performance input/output. This is critical for the prime goal of minimizing the time to obtain a solution.

Industry standards. Industry standards that easily permit the interconnection of multiple vendors in a heterogeneous environment are a key component of the TCS. Interconnection will be accomplished by providing the current widely used and evolving systems and protocols for IBM's technical computing products. These industry standards are found in both the systems and network layer.

Networking. The network layer, which provides interconnect and communications, is the primary glue that binds the UNIX operating environments to SAA and other vendor systems. There are two main aspects to communications. The first is the protocols or languages that the operating systems use to communicate with each other; the second is the distribution of data or function (e.g., programs, data files, screen/visualization frames, computation) to the elements networked together.

In technical computing, the protocols used by the operating systems to communicate with each other are based on TCP/IP. These will evolve to support the International Organization for Standardization (ISO) protocol suite. The TCP/IP suite of protocols is widely supported by many vendors and is available for IBM's key technical computing operating environments. The functions available to end users permit them to

Figure 1 Technical Computing Structure

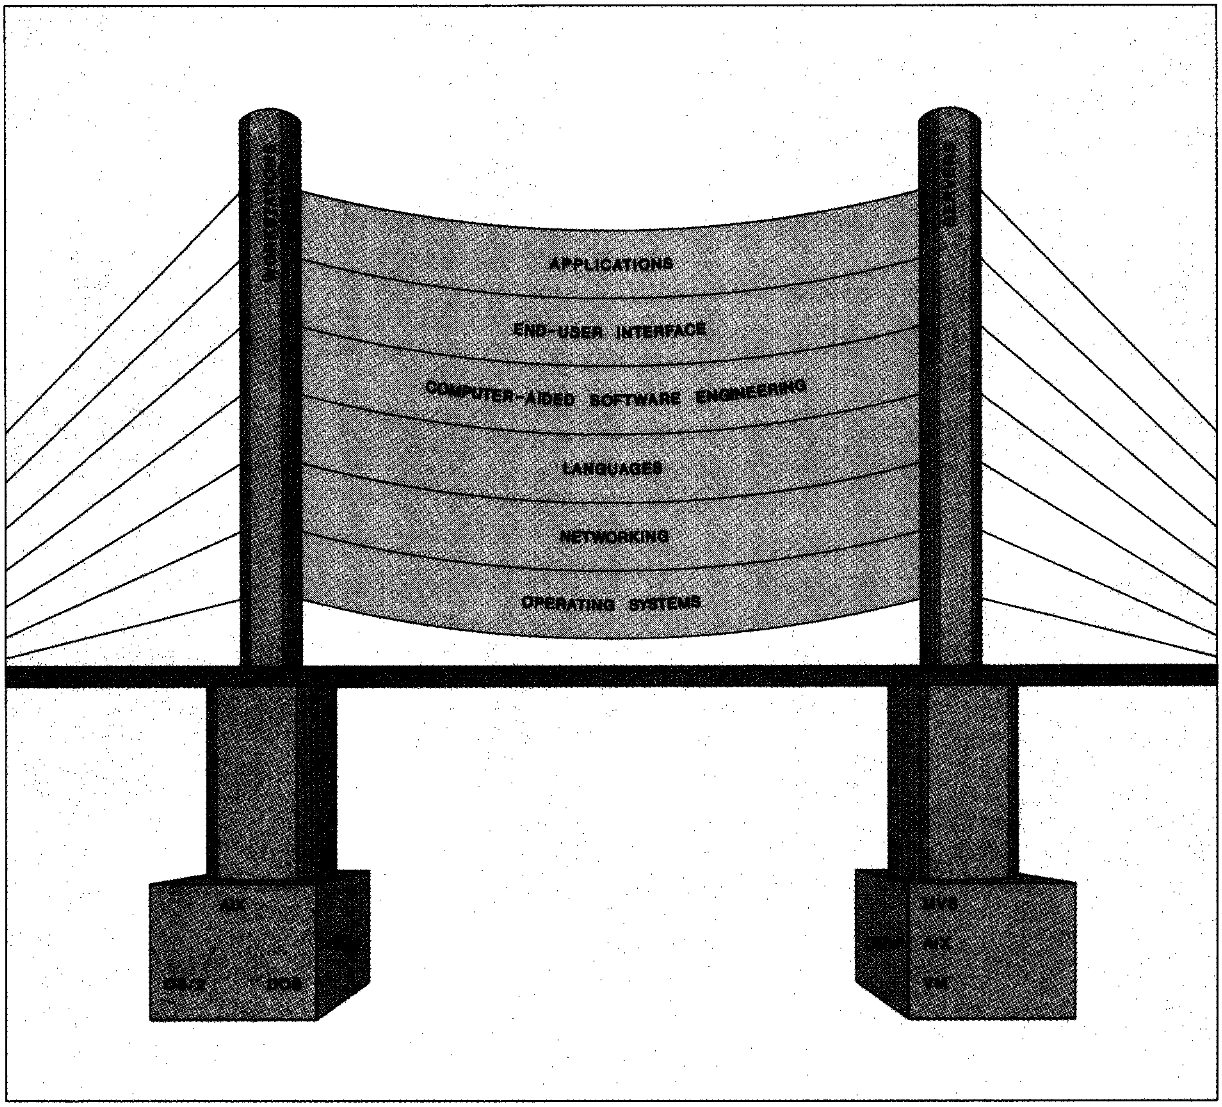

remotely access other computing resources (TELNET),[3] remotely execute commands, and transfer file and mail between computing resources (File Transfer Program, or FTP,[4] and SMTP).

The second aspect of communications is the distribution of data or function. One technique is X-Windows[5] client and server functions, supported by the IBM Advanced Interactive Executive™ (AIX™) workstation environments. Further, the IBM host environments, including AIX, Virtual Machine Facility (VM), and Multiple Virtual Storage (MVS) provide support for the X-Windows client. These functions provide the user with the ability to partition applications so that computation can be done on IBM supercomputers with graphics or visualization performed on IBM workstations.

For those environments where users need to share files, the Network File System (NFS®) provides functions that permit workstation users to transparently access host files and use IBM host systems such as AIX, VM, or MVS as file servers. Server implementations are also available on the RT®/System.

Another important tool, which permits users to distribute computation between systems, is the remote procedure call. This communication primitive permits users to call subroutines on different systems, thus segmenting their computations and assigning tasks to the computing resources that are best suited for a specific type of work. For example, while the user is interacting with a workstation, a large computation can be initiated in a supercomputer, and the results can then be manipulated and displayed on the workstation.

IBM has developed an extensive set of engineering and scientific subroutines (ESSL).

Systems and tools. The systems and tools component contains several key layers: systems, languages, CASE, and user interfaces. The AIX family is an important part of the TCS. With the AIX operating system, IBM made a major commitment to support and enhance UNIX operating systems for technical computing users. The AIX family incorporates most UNIX industry standards while providing additional capability. AIX is consistent with UNIX System V and selected Berkley Software Distribution functions. The added-value capabilities of AIX include: personal computer compatibility functions, breadth of product support across the Personal System/2® (PS/2®) Models 70 and 80, the RT/System, IBM 9370, 4381, and 3090 systems, and capabilities to bridge across UNIX systems as well as interconnect with IBM's SAAsystems.

The language layer includes compiler and run-time technology, object-oriented languages, and libraries. For instance, FORTRAN and C are standard across AIX and SAA, thus providing great advantages to the application developer. In areas where language extensions have become widely used in the industry, we have developed translation aids for the application developer to assist code migration. Our objectives with TCS languages, such as C, FORTRAN, and others that evolve, are to provide workstation and host affinity and portability between IBM and non-IBM systems, and to take advantage of the advances in hardware technology.

Another aspect included in the language layer is library support for the programming environment. IBM has developed an extensive set of engineering and scientific subroutines (ESSL) for common computational science functions. These subroutines provide tremendous productivity advantages while taking advantage of vector, parallel, and large memories in IBM's systems. Release 3 of ESSL provides parallel computational routines for two- and three-dimensional FFT (Fast Fourier Transform), matrix multiplication, and linear algebra. This is the first step toward more parallel algorithms and applications.

The Computer-Aided Software Engineering (CASE) layer of TCS will provide a full suite of tools to support the development life cycle. The characteristics of these tools include the heterogeneous environment of multiple vendors and recognition that the majority of application development is being done on workstations.

The end-user interface layer provides the workstation-based graphics and user interface. It includes window managers, application toolkits and dialog managers, and high-function graphics. Object-oriented techniques will become the dominant technology for the end-user interface.

IBM's approach, shown in Figure 2, builds on X-Windows Version X.11[5] with Display PostScript® support. In this environment, toolkits such as those developed at Carnegie Mellon University for the Andrew project and others can be used to further aid technical professional productivity. These toolkits, which run as application clients on X-Windows, are extensible object-oriented systems for the development, display, and manipulation of graphical and multi-media objects. Components such as text, vector graphics, raster images, equations, spreadsheets, and animation objects can exist both as standalone editor applications and as building blocks to be combined and embedded arbitrarily within one another.

Figure 2 End-user interface

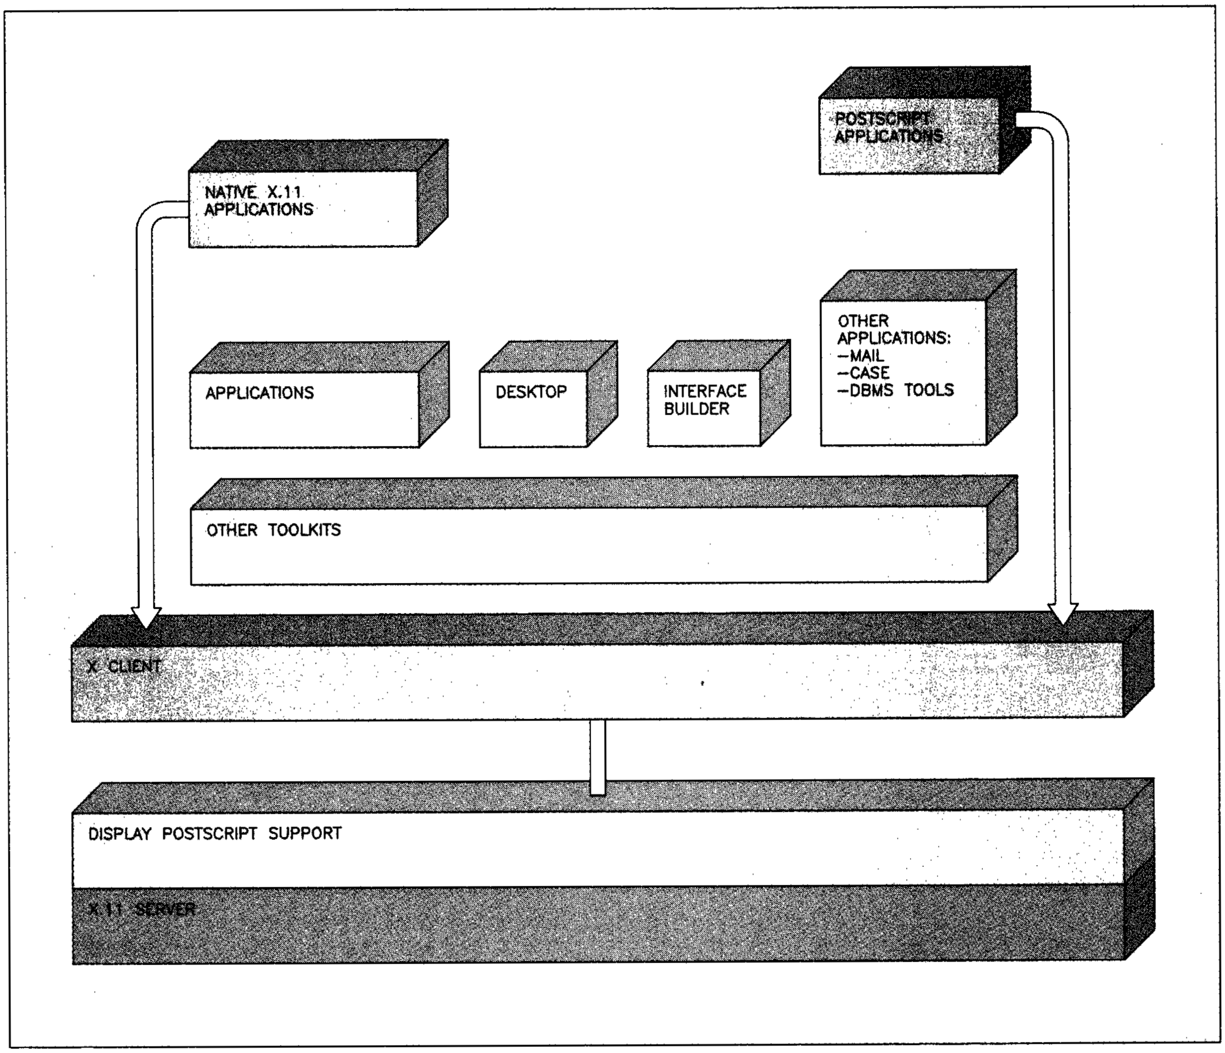

Applications. The applications layer uses the multiple dimensions of technical computing, such as vector, parallel, and visualization, to help minimize the technical computing user's time for obtaining solutions. This provides the ability to integrate the growing set of applications with the technical professional solution. These applications will often be software vendor applications that provide a specific function and can stand alone or be integrated across multiple offerings. These applications deal with the scientific or professional tasks that an end user performs.

An example of this layer for a *computational chemist*, a chemist who has to use computational facilities and select appropriate algorithms, is shown in Figure 3. The computational chemistry environment is comprised of several tasks in which the nature of the user's research, or the size of the scientific problem, determines the numerically intensive computing and graphic visualization requirements. Represented in Figure 3 is an overview of the activities associated with molecular modeling: model generation (chemical moiety information), molecular modeling (calculation), and model post-processing (interpretation). These activities may be distributed over a group of platforms—technical workstations (A and F), mainframe computers (C), or both (D). The technical workstation allows the user to graphically

Figure 3 Computational chemistry environment

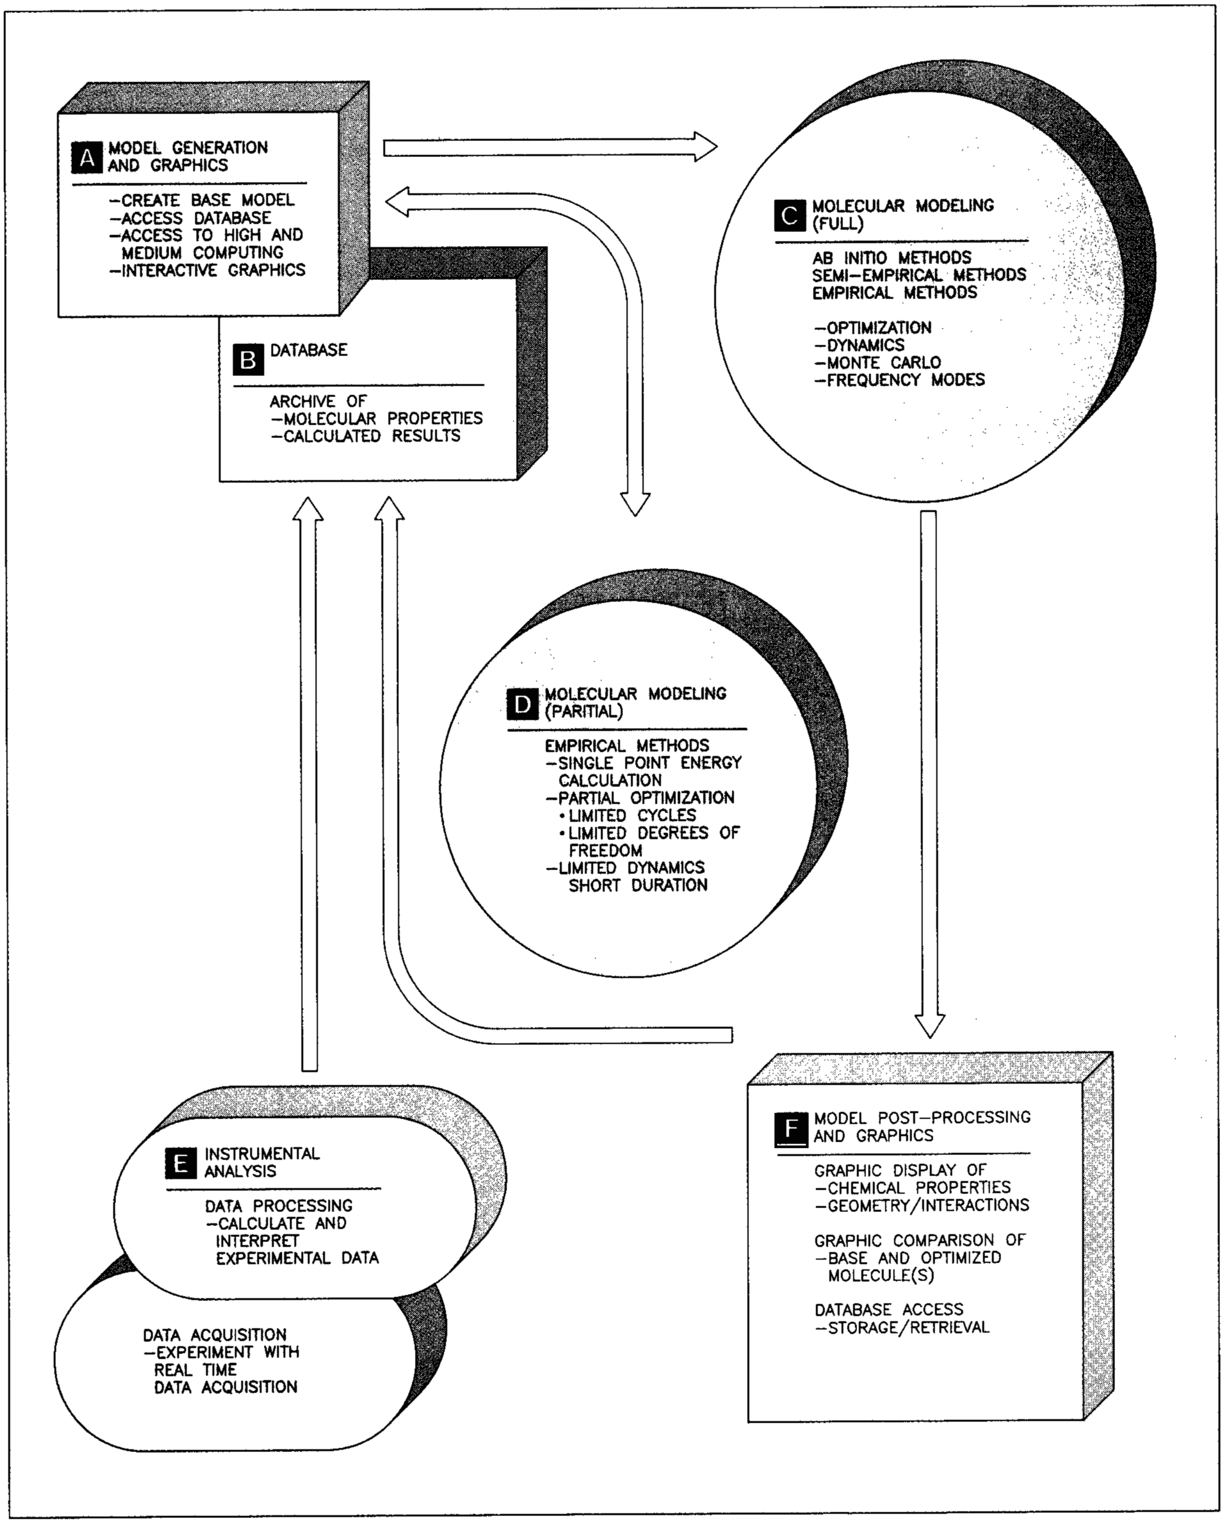

build and interpret chemical structures either derived from databases or generated as the result of a molecular modeling calculation. Depending upon the size

Technical computing is a primary tool of modern computational science and engineering.

of the chemical problem and the mathematical rigor desired, the molecular modeling calculation could be performed on a workstation, a mid-range system, or a supercomputer. Often preliminary calculations ("partial") on a workstation or mid-range system provide insights that may influence a larger calculation ("full") on the mainframe or supercomputer.

IBM's Computer Integrated Manufacturing, another example of this layer, allows the integration of CIM functions throughout an industrial enterprise. CIM builds upon the SAA and AIX structures by providing the foundation for a consistent base to store and manage a flow of both product and process data within the industrial organization.

Technical relationships

Technical computing is a primary tool of modern computational science and engineering. The functions and performance needed by the users to minimize their time to obtain a solution are often achieved through astute algorithm construction, careful usability analysis, and advances in technology. Through a better understanding of the entire process of problem solution, gained through relations with users, we can ensure that our product offerings evolve to better satisfy our customers.

These special business relationships are jointly funded projects with academia, government, and industry, and are an important part of our technical computing effort. They help us to better understand the systems and application enhancements that must be made to capitalize on new advances. Some of the work efforts are described below, and they are categorized into three areas of vital interest to IBM: (1) supercomputing and processor parallelism, (2) workstations and networks, and (3) UNIX and industry standards.

Supercomputing and parallelism. Cornell National Supercomputing Facility (CNSF) is one of five centers established by the National Science Foundation's Office of Advanced Scientific Computing in 1986 to provide supercomputer access to researchers. The Cornell center supports scientists and engineers throughout the country with two IBM 3090 Model 600Es with 12 Vector Facilities. Cornell has established a Smart Node Program with participation involving over 50 universities with 1600 users and 400 research projects. The Strategic User Program has 10 to 15 select projects that require huge amounts of resources and show exceptional promise for scientific breakthroughs.

One focus area for Cornell and IBM has been parallel processing using multiple independent processors on a single job on the Model 600E. This experience has shaped the recent announcement of IBM's Parallel FORTRAN Program, which automatically creates parallel programs to take advantage of multiple processors. In addition, the detailed analysis of end-user jobs has enabled IBM to determine performance improvements through hardware, compiler, and architectural changes. Finally, users are experimenting with new algorithms to improve applications operating in a parallel environment. Future focus areas for Cornell include parallelism across 12 processors and use of the AIX and UNIX systems on systems ranging from workstations to supercomputers.

IBM supports the European Academic Research Network and Supercomputer Centers to encourage supercomputer research, education, development, and exchange. The first of these supercomputer centers at the Centre National Universitaire Sud De Calcul in Montpellier, France, will be connected to the Cornell University National Science Foundation (NSF) supercomputing center by a high-speed (two megabits per second) link to the NSF network.

Los Alamos National Laboratory (LANL) and IBM have formed a special study in supercomputing to understand system requirements driven by the huge computational problems common in the national laboratory environments. LANL is one of the world's largest users of supercomputers. Historically, LANL computing requirements have been a precursor of the needs of other supercomputer users. Thus, a collaboration with LANL scientists provides IBM developers with an understanding of customer requirements in a prominent and demanding supercomputer center, and it will influence product directions. Particular areas of focus in this partnership include architecture definition, parallelism extensions, virtual memory studies for applications that were memory constrained on previous systems, and visualization.

This study will lead to improved algorithms for modeling the transonic flow about an aircraft.

One example of IBM's efforts is the joint study with Princeton University. This study will lead to improved algorithms for modeling the transonic flow about an aircraft to take advantage of the vector and parallel computing capabilities of the IBM 3090 family. Another example is the study with the University of California at Berkeley leading to better means to solve a variety of problems in integrated circuit design.

IBM and Supercomputing Systems, Inc. (SSI) have also formed a business partnership to develop a state-of-the-art high-end supercomputer. SSI is the supercomputer company formed by Dr. Steven Chen, formerly of Cray Research, Inc.

Workstations and networks. An example of a joint study designed to support the technical computing environment is the work at Syracuse University which prototypes a parallel software design environment for technical workstations. The initial study created a set of integrated tools to determine vectorization potential in FORTRAN programs and the associated performance improvement. Through extensions to the AIX command structure, a UNIX-like approach to submitting FORTRAN jobs to remote System/370 processors was also developed.

Current research using X-Windows extends the environment to the development of parallel FORTRAN programs. Parallel constructs are graphically displayed, performance effects of parallelization are analyzed, and optimization aids are presented.

Researchers across the country are benefiting from the new high-speed NSF network linking the supercomputing research centers and regional networks of the nation. IBM has developed and provided the IP routers as part of a team consisting of IBM, MCI, and a consortium of eight Michigan universities called MERIT, Inc. As outlined above, this network is managed by IBM's NetView, and the network will evolve over the next five years to the proposed OSI (Open System Interconnection) standards.

The UNIX system and standards. IBM is a participant in the standards development for POSIX[6] and X/Open®[7]; two recent actions reinforce our interest in current and emerging open standards.

IBM joined six other computer companies in announcing the creation of an international foundation to develop and provide a completely open software environment to make it easier for customers to use computers and software from many vendors. The Open Software Foundation (OSF) is dedicated to an open environment promoting portability, interoperability, and scalabilty. It will develop a software environment, including application interfaces, advanced system extensions, and a new operating system. The operating system will use core technology from a future version of IBM's AIX system as its development base.

IBM recently joined X-Consortium, a group of UNIX vendors that is enhancing in a compatible fashion the function of X-Windows, originally developed by the Massachusetts Institute of Technology's Project Athena.

These examples are only a few in an ongoing program covering many disciplines. As a result, the requirements of a wide variety of computational engineering and scientific disciplines will shape IBM's product developments to better solve customer requirements.

Conclusion

Technical users in academia and industry today are facing a new era of competitiveness. Their emphasis is on delivering higher quality products and services at lower costs, and the technical computing environment in which they work has become increasingly important in their drive to solve their problems in the shortest possible time.

The purpose of the Technical Computing Structure is to provide a framework for the technical users to achieve their goals, and for an enterprise to manage a system for the technical computing users. This Technical Computing Structure will also help IBM establish the right programs to deliver the vision of a seamless and efficient technical computing environment.

Acknowledgments

The authors thank J. B. Archer, D. D. Cerutti, H. Brauer, J. L. Latta, and R. F. Wolfe for their many valuable discussions on content and form. In addition, many other people working at multiple IBM sites (Kingston, Santa Teresa, Austin, Milford, Palo Alto, Dallas, Yorktown, White Plains) have provided insight into the material discussed in this article.

UNIX is developed and licensed by AT&T and is a registered trademark of AT&T in the U.S.A. and other countries.

NetView, PROFS, Advanced Interactive Executive, and AIX are trademarks, and Personal System/2, PS/2, and RT are registered trademarks of International Business Machines Corporation.

PostScript is a registered trademark of Adobe, Inc.

NFS is a registered trademark of SUN Microsystems, Inc.

X/Open is a registered trademark of X/Open Co., Ltd.

Cited notes

1. TCP/IP is a set of layered communications protocols defined by the Department of Defense in the 1970s for multivendor connectivity for local and wide-area networks. It is being used by the National Science Foundation as the protocol to link the six supercomputer centers and seven mid-level academic networks across the country.
2. SMTP is a standard function provided within TCP/IP for transferring electronic mail among the units on the network.
3. TELNET provides access to terminal emulation communications applications as a function within TCP/IP.
4. FTP is another standard function in TCP/IP for support of file transfers among the systems and workstations in the network.
5. X-Windows and X-Windows Version X.11 are defined, public domain functions provided within TCP/IP to manage windows for the end user.
6. POSIX is an IEEE standard for an interface standard on UNIX systems.
7. X/OPEN is an organization of users (e.g., Aetna Life & Casualty Corp., Bankamerica Corp., and others), hardware vendors (e.g., DEC, AT&T, and others), and software vendors (e.g., Oracle Corp., Informix Software, Inc., and others); committed to creating an open standards universe based on the UNIX platform.

General references

E. F. Wheeler and A. G. Ganek, "Introduction to Systems Application Architecture," *IBM Systems Journal* **27**, No. 3, 250–263 (1988).

V. Ahuja, "Common Communications Support in Systems Application Architecture," *IBM Systems Journal* **27**, No. 3, 264–280 (1988).

D. E. Wolford, "Application enabling in SAA," *IBM Systems Journal* **27**, No. 3, 301–305 (1988).

L. A. Buchwald, R. W. Davison, and W. P. Stevens, "Integrating applications with SAA," *IBM Systems Journal* **27**, No. 3, 315–324 (1988).

W. P. Dunfee, J. D. McGehe, R. C. Rauf, and K. O. Shipp, "Designing SAA applications and user interfaces," *IBM Systems Journal* **27**, No. 3, 325–347 (1988).

R. A. Demers, "Distributed files for SAA," *IBM Systems Journal* **27**, No. 3, 348–361 (1988).

R. Reinsch, "Distributed database for SAA," *IBM Systems Journal* **27**, No. 3, 362–369 (1988).

A. L. Scherr, "SAA distributed processing," *IBM Systems Journal* **27**, No. 3, 370–383 (1988).

L. K. Loucks and C. H. Sauer, "Advanced Interactive Executive (AIX) operating system overview," *IBM Systems Journal* **26**, No. 4, 326–345 (1987).

R. Q. Cordell II, M. Misra, and R. F. Wolfe, "Advanced Interactive Executive program development environment," *IBM Systems Journal* **26**, No. 4, 361–382 (1987).

S. G. Tucker, "The IBM 3090 system: An overview," *IBM Systems Journal* **25**, No. 1, 4–19 (1986).

Y. Singh, G. M. King, and J. W. Anderson, "IBM 3090 performance: A balanced system approach," *IBM Systems Journal* **25**, No. 1, 20–35 (1986).

D. H. Gibson, D. W. Rain, and H. W. Walsh, "Engineering and scientific processing on the IBM 3090," *IBM Systems Journal* **25**, No. 1, 36–50 (1986).

W. Buchholz, "The IBM System/370 vector architecture," *IBM Systems Journal* **25**, No. 1, 51–62 (1986).

R. S. Clark and T. L. Wilson, "Vector system performance of the IBM 3090," *IBM Systems Journal* **25**, No. 1, 63–82 (1986).

J. K. Kravitz, D. Lieber, F. H. Robbins, and J. M. Palermo, "Workstations and mainframe computers working together," *IBM Systems Journal* **25**, No. 1, 116–128 (1986).

R. C. Agarwal, J. W. Cooley, F. G. Gustavson, J. B. Shearer, G. Slishman, and B. Tuckerman, "New scalar and vector elementary functions for the IBM System/370," *IBM Journal of Research and Development* **30**, No. 2, 126–144 (March 1986).

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

W. S. Humphrey, "The IBM large-systems software development process: Objectives and direction," *IBM Systems Journal* **24**, No. 2, 76–78 (1985).

R. A. Radice, N. K. Roth, A. C. O'Hara, Jr., and W. A. Ciarfella, "A programming process architecture," *IBM Systems Journal* **24**, No. 2, 79–90 (1985).

Patti L. Prairie *IBM Data Systems Division, P.O. Box 100, Kingston, New York 12401.* Ms. Prairie is Director, Technical Computing Systems and Applications, in the Data Systems Division. She joined IBM in 1973 at the Rochester, Minnesota, laboratory. Her responsibilities there included serving as manager of Rochester Design Systems in 1981, focusing on engineering design systems and hardware technology developments. She began a site software development tools group in 1982, which was incorporated into division-wide product development support. During 1983 she was the technical assistant to the System Products Division vice president of development. In 1984 she was manager of development productivity at SPD headquarters. Subsequently, she was the System/88 product manager responsible for hardware and software development. Ms. Prairie joined the DSD organization in 1986 and was responsible for IBM's development strategy, related planning activities, and market development for IBM's numerically intensive computing hardware and software, including the 3090 Vector Facility. In April 1988 she was appointed to her current position where she is responsible for ensuring that IBM provides hardware, software, and system products to satisfy the needs of the technical computing professional. She has also been active in industry and university relations. She participated in the preparation of design automation courses at the University of Minnesota and taught electrical engineering courses in the graduate school there. Ms. Prairie graduated with high honors from the University of Minnesota with a B.S. in mathematics and economics and an M.S. in electrical engineering and computer science.

Allan H. Weis *IBM Data Systems Division, 44 South Broadway, White Plains, New York 10601.* Mr. Weis is Vice President, Data Systems Division, IBM Enterprise Systems. Mr. Weis joined IBM in 1961 as a trainee in the New York Service Bureau. In 1968 he moved to the Research Division where he managed several advanced technology projects in communications and data management. He was assigned to the Zurich laboratory in 1972 as assistant to the director of the laboratory, and two years later he was appointed manager of the European Laboratory Computation Network in Hursley, England. Upon his return to the United States, he was appointed to the Corporate Engineering Programming and Technology staff, and became acting director of systems and programming. In 1980 he joined the Systems Communication Division where he held several management positions. In 1983 Mr. Weis was appointed director of computing systems in the Research Division where he was responsible for the computing and communications systems and related advanced technology programs across IBM's Research laboratories. In April 1988 he was appointed DSD vice president, engineering/scientific computing, with worldwide responsibility for the strategy, development, and technical support of IBM's large systems for numerically intensive computing. Mr. Weis majored in electrical engineering at the University of Kansas, and received his Science Masters from the Massachusetts Institute of Technology on an Alfred P. Sloan Fellowship. He is a member of the National Research Council Committee on the use of Information Technology in Research, a member of the board of directors of the New York State Education Research Network (NYSERNET), a member of the College of Engineering Industrial Advisory Board at the University of Illinois at Chicago, and has served on the National Science Foundation Network Executive Committee.

Reprint Order No. G321-5334.