Final Report

An Integrated Oncology Workstation

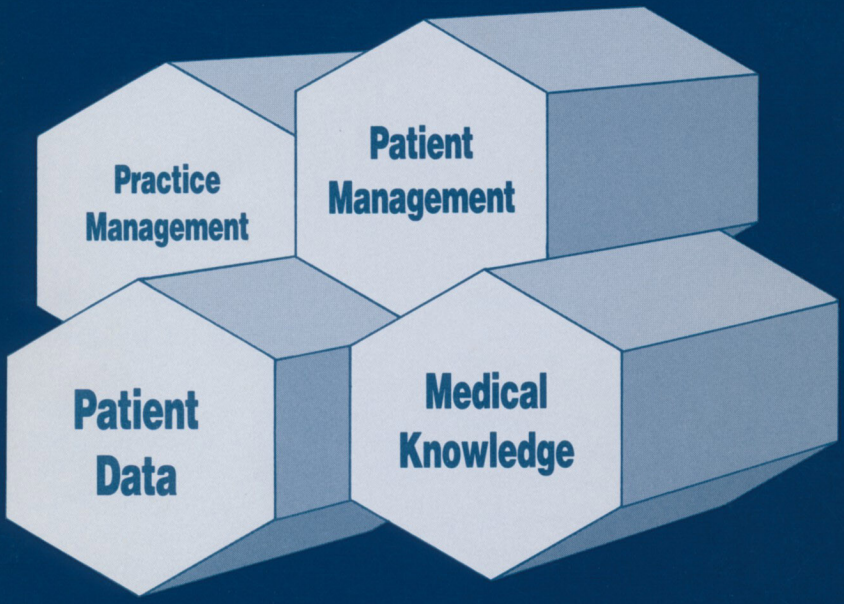

Final Report

An Integrated Oncology Workstation

Submitted Pursuant to
Contract NIH 263-89-C-0080

Submitted to:
Robert J. Esterhay, Jr., M.D.
Project Officer
May 1990

National Cancer Institute, NIH
Building 82, Room 201
9000 Rockville Pike
Bethesda, Maryland 20892

Submitted by:
Edward H. Shortliffe, M.D., Ph.D.
Principal Investigator
Clifford E. Wulfman, M.S.*
Thomas C. Rindfleisch*
Robert W. Carlson, M.D.†

Department of Medicine
* Section on Medical Informatics (General Internal Medicine)
† Division of Oncology
Stanford University School of Medicine
Stanford, California 94305-5479

Table of Contents

Appendices

Foreword

INFORMATION TECHNOLOGY today has expanded far beyond the mainframe computers of a few years ago to include powerful yet affordable desk-top personal computers. Such technical workstations can include high-speed local and remote communication networks, image processing (including facsimile transmission), expert systems, and interactive voice systems. These new technologies are rapidly maturing and are beginning to be utilized individually in practicing physicians' offices, but not yet in a broad, integrated fashion for the documentation and management of patient care and clinical research.

In a clinical research setting at Stanford University, Dr. Ted Shortliffe and his group have demonstrated that computer software can assist the physician in determining patient eligibility and in enrolling patients on clinical trials. They have also demonstrated, with the ONCOCIN research project, that expert system software can support the physician in the proper application of the test and treatment details for a patient enrolled on a trial and can document all patient care and clinical trial activities.

In 1983 the National Cancer Institute (NCI) initiated the Community Clinical Oncology Program (CCOP) to increase the involvement of community physicians in clinical trial research on cancer treatment, prevention, and control. By involving more community-based physicians in the clinical trial process, either through their increased referrals to research centers or by their personal involvement as investigators, the CCOP initiative has sought to assure that patient accrual to clinical trials will increase. Today, CCOP and other community-based initiatives such as the Cooperative Group Outreach Program are contributing to patient accrual to the point where more than 50% of all patients entered in formal NCI clinical research protocols are now entered by community oncologists.

It is to this group of community-based physicians that the *integrated oncology workstation* will be initially directed. NCI appreciates the required time and logistical impediments to participation in the clinical trial process by community physicians. The *integrated oncology workstation* is envisioned as the ultimate solution for assisting the physician and his or her office staff in recording data, keeping records, planning and managing patient care, and in accessing medical knowledge.

The conceptual overview of the *integrated oncology workstation*, provided in this monograph, was written to describe what the features of this clinical care system for practicing oncologists might include. The clinical workstation for oncology will be organized around a database that contains the current information for each patient being followed by the physician (the "electronic medical record"). It will also provide facile access to a knowledge base of information resources such as PDQ, CANCERLIT and MEDLINE. This concept — the integration of data and knowledge — is the foundation on which the clinical workstation for oncology will be built.

At the request of the NCI, Dr. Shortliffe and his colleagues have described a provocative vision of the integrated workstation for oncology. Their analysis has included policy and planning issues and has related what is possible to an assessment of state-of-the-art computer technology. We hope that this monograph will generate in practicing oncologists, academia, industry, and the government an awareness and interest that will advance the commitment to create and make available an *integrated oncology workstation.*

Robert J. Esterhay, Jr., M.D.

Section I
Executive Summary

As clinical practice struggles to deal with an increasingly complex set of administrative, reporting, and fiscal requirements, it is small wonder that many practitioners have found it difficult to become active participants in clinical research. Despite active efforts to encourage increased participation in cancer clinical trials, only a fraction of oncologists in the United States regularly enroll patients in the very trials which hold the hope of the future for patients with cancers. We believe that computing technology offers solutions, both for dealing with the current chaos in data management, paperwork, and reporting required of all practitioners *and* for facilitating the participation of physicians in formal clinical trials. This report, prepared at the request of the National Cancer Institute, summarizes that technology and seeks to provide a vision of the future that is at once both exciting and practical. Our goal is to provide a conceptual overview of what the integrated oncologist's workstation of tomorrow could and should be. By an *integrated oncology workstation* we mean an advanced personal computer that provides an oncologist with a wide variety of data-management and decision-support tools, many of which run on other computers but which are accessed via networks or telephone/modem connections. The notion of integration requires that the various resources are tied together in a uniform fashion which allows the user to move fluidly from one application to another, generally using the same interactive methods, without needing to learn a variety of different conventions depending upon which program is being used at any given time.

The striking lesson of the report is that *most of the barriers to successful implementation of a fully integrated oncology workstation are not technical but logistical and sociopolitical.* Although there is always room for refinement of methodologies and for new research at the frontiers of computer science, the technical foundations for what we propose already

exist and are being used effectively today in areas outside of medicine. Even within medicine, demonstration prototypes exist. They have shown that computers can assist oncologists and other health workers in a number of ways:

- ***As medical record keeping tools:*** Computers can store very large volumes of data in a variety of formats, including numeric data, text, and images, and they can present those data in ways tailored to the particular needs of individual users. Computers can therefore be used to maintain patient records in a way that makes them simultaneously available to different healthcare providers.
- ***As gateways to medical knowledge:*** Many sources of medical information can now be accessed with computers, including bibliographic databases, databases of clinical trial information, and medical textbooks.
- ***As sources of medical decision support:*** Programs can assist physicians and other health workers in a variety of clinical tasks, such as diagnosis and protocol management.

The 1980s saw many rapid advances in the evolution of computer technology to support the development and dissemination of ever larger and more sophisticated systems:

- Computer hardware became much faster and much cheaper
- The capacity to store data and retrieve it quickly increased enormously
- Graphical interfaces made computers much easier to use
- New software techniques made it easier to use large computerized information bases

These advances enabled the development of a number of computer-based tools specifically for clinical oncologists. Among these were:

- The Physician Data Query system (PDQ), developed by the International Cancer Information Center (ICIC) of the National Cancer Institute (NCI), designed to provide practicing physicians with the most current information available on cancer treatment options and available clinical research trials;
- OCIS, the Oncology Clinical Information System which coordinates all data-management activities in the oncology clinic at Johns Hopkins Hospital;
- ONCOCIN, our own large decision-support system designed to assist clinical oncologists with the management of patients

receiving protocol-based chemotherapy;

- CANCERLIT, a bibliographic database produced by the NCI;
- The complete texts of important oncologic reference works, stored with PDQ and CANCERLIT on high-capacity digital optical discs (CD-ROMs), in a product known as OncoDisc.

Despite these successes, experience in the 1980s showed that it is difficult to attract clinical users to such tools when they are not well integrated with their routine patient-care tasks. As our understanding of integrated workstation capabilities and networked linkages has matured, a new model of the potential for computers in clinical oncology has emerged.

It is now clear that if the full power of each of these new tools is to be made widely available, the tools must be brought together in a single oncologist's workstation, where their combined utility can create the critical mass of functionality needed to draw the computer into clinicians' routine of clinical care. Furthermore, the tools must be unified by natural human-computer interfaces that are intuitive to use and that require minimal training.

In Section II of this report, we describe an idealized integrated workstation for clinical oncology, but we fully recognize that many of the individual components will be challenging to implement and to deliver in a cost-effective manner. Such an environment will not be developed overnight, nor will it be possible to introduce the system into clinical settings in a single step. Although it is beyond the scope of this document to provide a detailed design specification or a timeline for development, our report presents some of the strategies that might be used to implement an integrated oncologist's workstation and some of the issues that must be addressed during its design and development.

- The workstation's developers should adhere closely to standards in their choice of basic systems support and in the design of their own modules, thereby making it more likely that the oncologist's workstation could be deployed on a variety of computers from different manufacturers.
- The need for all the components of the oncologist's workstation to work as an integrated whole should be of primary concern in the early design stages.

- The disseminators of an oncologist's workstation should not plan to reach all of its broad target community at once but should instead identify an initial, smaller subcommunity, one characterized by its eagerness to acquire new technology and by the relative ease with which practical and technological barriers to the integration of new technology can be overcome. We performed a survey of oncologists which suggests that small-to-moderately–sized private oncology clinics constitute the most promising initial market because they tend to be more autonomous and more anxious to integrate computer technology into their practices.
- The full functionality of an integrated oncologist's workstation should unfold in a series of staged releases that correspond to the achievement of major technical milestones that support the features in each of the broad usage categories outlined in Section II.E. We suggest, in Section II.F, one approach to identifying staged technical milestones and features.

The workstation will be installed in clinical and oncology research settings that will vary considerably in their size, sophistication, and technical resources. Its design should therefore be flexible enough to accommodate site variability without modification of its intrinsic components. As much as possible, the workstation's configuration should be under the control of software "switches" that can allow the system to turn off unneeded functionality (and turn it on again if a site's requirements change) and to establish the appropriate pathways for information access and data exchange.

Impediments to the Vision

We have been careful in this report to envision an oncologist's workstation that could be implemented with technology already available today. Nevertheless, there are many nontechnical obstacles to the development and deployment of such a system. The removal of these obstacles is the responsibility not of individual developers, who can only react to what has been put in place, but of planners and policy makers. The impediments to the vision of an integrated oncologist's workstation are discussed in Section III and summarized here.

Attitudes of Health Workers

Although physicians and other health workers have shown increasing interest in computing issues, they express recurring concerns when asked

to consider using computers in their daily patient-care work. These attitudes can constitute significant barriers and, accordingly, should be addressed both in the system's design and in the educational and training materials that are produced. Typical issues discussed fully in Section III include:

- Fear of loss of rapport
- Fear of loss of control
- Inertia
- Fear of active decision support
- Fear of legal liability
- The challenge of data entry
- Reliance on the younger generation

Costs of Automation

Although outright costs of equipment and software may not be trivial, there is ample evidence that physicians and group practices will invest heavily in technology when it is perceived as offering them clear advantages in terms of efficiency, quality of care, or the ability to attract patients. However, system developers must address other costs, particularly the start-up expenses associated with change. These costs are often more than monetary; fears of disruption, dislocation, and loss of autonomy are often associated with the introduction of computer technology into existing patterns of work. Carefully designed phase-in strategies and user-education programs are therefore crucial if transition expenses are to be minimized or, at least, spread out over many months.

Lack of Standards and Planning

An impediment that is poorly appreciated by end users but that places a key constraint on system developers is the lack of established standards for data sharing, terminology, and computer-to-computer communication. The establishment of standard data dictionaries for medical information systems is crucial if diverse machines and practices are to share information. The need for a common network infrastructure is particularly critical. The oncology workstation model we have proposed depends, for its optimal realization, on a coordinated plan for local and nationwide connectivity of computers. To establish such an infrastructure for biomedical computing requires planning at the community, regional, and national levels.

Technology Transfer Issues

An obvious question that arises when one considers the oncology workstation model proposed in this document is "Who is going to build and market it?" Industry perceives the risks in bringing a large software product to the medical market to be high, and many past failures discourage even the most adventurous. Strategies for dissemination that involve shared investment and diluted risk, possibly with government assistance or coordination, are likely to be necessary.

Trends in Cancer Clinical Trials and Their Coordination

As our knowledge of cancer treatment has increased, cancer clinical trials have become increasingly complex and frequently require the collaborative effort and cooperation of large numbers of geographically dispersed physicians and patients. Because large regional and national organizations have been formed to coordinate the implementation of clinical trials, the resulting need for close cooperation and coordination has become both a mandate for enhanced networking with computer-based data capture and, ironically, an impediment to the effective adoption of new technologies. Independent clinics or medical centers cannot make independent decisions regarding the use of electronic databases and computer-based reporting methods; they must look to the central organizations for direction.

Conclusions

In this report, we hope to have conveyed our excitement about what is possible, tempered with a realistic sense of the significant barriers to effective and timely implementation of the vision we have described. The reader who develops an enthusiasm for the notions embodied in this report, and who sees the remarkably positive effect that such technology could have both on oncology practice and on clinical research, may well ask "How much will this cost?". Although individual workstations with the capabilities we have described are already available in the range of $10,000 per unit, the costs for software development and support, as well as for the networking and communications infrastructure, are much more difficult to predict. Amortized over the entire community, these costs are likely to create only a small incremental expense for the individual user, so it is the initiation and coordination of such efforts which are the

principal constraints on rapid progress in the field. As we emphasize in Section III, effective national leadership and standards setting, coupled with subsidies for creation of the communications infrastructure, will be crucial elements to enable the effective implementation of integrated clinical workstations. We must look to local, regional, and national health planners for the initiation of such programs. Only then will researchers and vendors be able to build systems that integrate well with the health-care environment.

Section II

The Integrated Workstation for Oncology

Contents

A. Background and History

For several years, our group at Stanford University School of Medicine has been engaged in the development and testing of an oncology protocol management system known as ONCOCIN. The program is a large system designed to assist clinical oncologists with the management of patients receiving chemotherapy regimens. It addresses issues of proper dosing and appropriate response to developing toxicities while helping to ensure the completeness and consistency of data that are collected for patients enrolled in formal clinical trials. The program began as a tool that ran on large time-shared computers using simple computer terminals installed in the Stanford Oncology Day Care Center. Later ONCOCIN was rewritten to run on single-user workstations that combined advanced programming tools with high-quality graphical display capabilities. ONCOCIN has been used on a limited basis by faculty and fellows at Stanford, and it has been formally evaluated in trials designed to assess its impact on data collection and its ability to provide appropriate management advice.

During the decade of this work, several other computer-based tools for clinical oncologists have been developed at other institutions. Particularly noteworthy has been the Physician Data Query system (PDQ), developed by workers at the International Cancer Information Center (ICIC) of the NCI. PDQ is designed to provide practicing physicians with the most current information available on cancer treatment options and available clinical research trials. Its database is meticulously maintained with current information, and the information is available both by telephone access to computers at the National Library of Medicine and via CD-ROM versions to which individuals or institutions may subscribe. The NCI also produces a number of other computer-based information resources, including a bibliographic database known as CANCERLIT.

One lesson of the last decade's experience has been the difficulty in attracting clinical users to computer-based tools when those tools are not well integrated with their routine patient-care tasks. As a research prototype, ONCOCIN has been used for only a fraction of the patients under treatment in the Stanford clinic, and its use has accordingly never been a natural part of the routine. Similarly, PDQ requires incremental effort if a physician is to access it, whether he or she uses the modem/network or

the CD-ROM versions. A natural inertia, coupled with unfamiliarity with computing technology, has limited both systems' use in ways that are discouraging to those who know the clinical value of the information that they have to offer.

These experiences have made it clear that despite their usefulness, these powerful new tools for clinical oncology will not be appreciated by their intended users if they are distributed piecemeal. As our understanding of integrated workstation capabilities and networked linkages has matured, a new model of the potential for computers in clinical oncology has emerged. What is required is an integration of this wide array of tools into a single computer-based workstation for clinical oncology. Only such an integration of functionality, coupled with natural human-computer interfaces that are intuitive to use and that require minimal training, will be able to achieve the critical mass needed to make computer-based tools part of the oncologist's clinical routine.

In this report we are responding to a request from the National Cancer Institute that we provide a conceptual overview of what the integrated oncologist's workstation of tomorrow should be. It will quickly become clear that such a workstation should not be a single consistent entity across all possible users and uses. Furthermore, the ideal tool cannot be viewed as an abstract futuristic notion but needs to be considered in the context of a strategy for the phased introduction of key components and related infrastructure over time.

The report is organized into three sections with three appendices. An introductory executive summary is intended for non-technical readers (principally physicians and health managers) who want to know the key issues in the design and development of an integrated workstation for clinical oncology. In the current section we provide a detailed, non-technical overview of the integrated workstation notion. Sections II.B–II.F provide a description of the oncologist's needs in an integrated workstation for oncology and of the methods for fulfilling these needs. Specifically, Section II.B reports on a clinical survey we performed to assess attitudes and needs of clinicians and other workers in a variety of oncology practice settings. The results of this analysis help provide a framework for considering the key functionalities that will need to be included in an integrated workstation for use by oncologists. Section II.C then provides a brief survey of modern computing technology, defining in

simple terms the key topics that need to be explained before the development and logistics of the oncologist's workstation can be understood. Section II.D provides three scenarios that describe some of the proposed alternate uses of the oncology workstation: use by the oncologist in private practice, by a clinical-trial data manager, and by the research director of a group overseeing cancer chemotherapy trials. Building on this background, we then provide a detailed discussion of the oncologist's workstation in Section II.E. Strategies for staged implementation of the integrated tools are then outlined in Section II.F.

Section III takes a step back to consider some of the policy and planning issues that will be crucial catalysts to the development of the workstation capabilities we are proposing. These include the need for local, regional, and national infrastructures for planning and coordinating data sharing and communications, as well as solutions to the individual practitioner's need for a stable industry that understands current practice and referral patterns and that can provide responsive support at a reasonable cost for the computer-based tools that it markets.

For readers who are especially interested in technical details of the proposal, two appendices provide a brief assessment of the state of the art in computer technology (Appendix A) and a discussion of pertinent strategies for selecting the appropriate hardware and software for a development task such as the one we are outlining (Appendix B). The discussion in Section II, on the other hand, specifically avoids any discussion of specific computers, manufacturers, operating systems, or programming languages.

Finally, we have developed a demonstration prototype that helps illustrate the kinds of features that should be possible when the integrated oncologist's workstation is developed. The sample screens in Section II.D were derived from this software package.

B. Report on a Clinical Survey

At the outset of our research for this report, we conducted a survey of a variety of outpatient oncology clinics to ground our design of an integrated oncologist's workstation on the broad needs of typical oncology practices. Although we have worked closely with the outpatient oncology clinic at the Stanford University Medical Center for many years, we wanted to be cautious in generalizing our experience and conclusions from that work to other settings. We therefore prepared a questionnaire and interviewed physicians and other personnel at five clinics within a radius of 100 miles of the San Francisco Bay Area. The types of clinics covered a wide spectrum:

- a 2-person private clinic that practices oncology part-time
- a 2-person private clinic that practices oncology full-time
- a large, full-time private oncology clinic
- a large oncology clinic associated with a university other than Stanford
- a 1-person practice within a large health-maintenance organization

The following section summarizes the features and needs the clinics held in common and the ways in which they differed. We discuss the nature of the practices we visited and the implications of our observations for the overall design of the integrated oncologist's workstation: the fundamental differences that make a single monolithic design solution impossible, and the principal barriers we identified that could impede the introduction and acceptance of an integrated oncologist's workstation.

B.1. The Nature of the Surveyed Practices

We tried to choose clinics for our survey that would show us as wide a range of practice sizes, settings, and styles as possible. One of us (R.W.C.) is a Stanford oncologist who was able to identify appropriate clinics and to make the required introductions. In all but one of the clinics we approached, the physicians were more than willing to talk with us; most spent several hours in the interviews. Our survey questions focused on clinic demographics, the amount of computer support the

clinic already received, areas in which the clinic already believed improved computer support could or could not be helpful, and the attitudes of the practice members toward the use of computers in their offices or clinics.

1.1. Clinic Demographics

As outlined above, the clinics ranged in size from a 2-physician private practice with 3–4 nurses and a few support personnel to a university-affiliated clinic with 6 full-time oncologists, 3 fellows, 7 data managers, and numerous other clerical and nursing staff. The apportionment of tasks varied considerably, sometimes out of preference, sometimes out of necessity. One of the widest disparities in task apportionment occurred in the allocation of clerical duties to physicians. In one clinic, physicians were filling out laboratory forms, admission forms, insurance forms, and many other documents that apparently could have been completed by a clerical staff member. In another clinic, one of the physicians voluntarily acted as data manager for the numerous protocols in which the clinic participated.

The practices' patient populations ranged widely, from 13 patients a day to over 40 patients per day. Our questionnaire did not inquire about the number of new patients, but it appears that for many of the clinics we surveyed, approximately 7% of the patients seen in a day were new patients.

Participation in clinical trials ranged from less than 5% of patients to over 15% of patients. The physicians in clinics with low clinical-trial participation cited administrative overhead as the chief barrier to increased involvement. They believed a computer system that helped with eligibility screening and data management could significantly increase the number of patients they could enroll on protocols.

1.2. Amount of Existing Computer Support

Every clinic we visited had a computer in it, but the similarity ended there. Some clinics had a few stand-alone PCs that were used for word processing by clerical and data-management staff (and sometimes by physicians); in one of the clinics, each physician with a computer had a different brand; in another clinic, one of the physicians had built his own patient-record system; still others had made attempts (some abortive) to

install or maintain commercial systems that provided (or purported to provide) integrated solutions to scheduling, billing, and clinical data management. None of the clinics had connections to networks. Some physicians had modems, either at home or in their offices, which they used to access MEDLINE or PDQ, but use of online services was idiosyncratic and highly dependent upon local computer lore. Most had heard of PDQ and expressed some interest in gaining access to it, often in self-defense (because their patients were arriving with PDQ printouts that they had obtained prior to their visits) or, occasionally, because they feared potential legal liability if they failed to consult the system.

1.3. Attitude Towards Clinical Use of Computers

We found among those we interviewed very few who had a realistic perception of the potentials for computers in their practices. Those who were computer-naive were sometimes highly enthusiastic about automation and unrealistically optimistic about the possibilities for functionality and integration. (They were unaware of the complexities of networking machines, protocols for data exchange among diverse hardware and software systems, systems-support requirements, and the like; to them, all these systems were computers, and computers should be able to communicate with one another.) Those who had worked with personal computers were often bound by their limited experience; they frequently found it difficult to envision hardware or software more powerful than that with which they were familiar, and they were accordingly skeptical that an integrated oncologist's workstation was achievable. Those who had been exposed to experimental systems in universities (either during their training or in the course of collaboration with medical informatics researchers) tended to have a more clear-eyed view of the possibilities and the limitations of future automation, as did those who had had extensive experience with failed office-automation solutions.

In one clinic, we had the opportunity to interview several data managers who were quite skeptical about the benefits a computer system could provide. Although they had enrolled 60 new patients on one research group's protocols in the previous 8 months alone, they did not think a computer would help in managing all these protocol patients. One pointed out that most study groups have their own sets of forms and insist that data be transmitted on them. If the study groups will not accept

computer-based forms, the data manager said, automating data management would simply make more work. She felt that an online patient calendar for tracking the scheduling of protocol-directed laboratory tests could be useful, but only if it were linked to the laboratory in the hospital. When we suggested a computer system might aid them in screening patients for protocol eligibility, the data managers replied that they did not see that a computer system would give them anything they did not already have. They had a binder filled with hard copies of protocol documents, grouped by disease and various other criteria, with which they felt they could very quickly determine a patient's eligibility for protocols.

We attribute the data managers' skepticism about computer support to two factors. The first is a relative unfamiliarity with the power of modern computing technology; they were unaware, for example, that off-the-shelf software can be used to produce computer-generated forms that are precise facsimiles of existing printed forms. The second is the stress induced by working near capacity. Because they were only just staying ahead of the paperwork they had to do, the data managers were unwilling to speculate about tinkering with their data-management system that, while not perfect, was at least functioning.

Their concerns, and similar ones pointed out by nurses and office support staff, illustrate the need to address logistical issues as well as large-scale coordination in the design of the integrated oncologist's workstation. A project to design and implement such a workstation must consult with data-collection agencies, hospital-information-system and laboratory-machine vendors, insurance agencies, and other entities that require access to clinical data. Where coordination is impossible, the workstation must be designed with sufficient flexibility to enable individual sites to tailor their installations to import and export clinical data in the formats they require.

B.2. Implications for the Workstation Design

2.1. Common Threads

The overriding feature shared by all the clinics we visited was their desire to have some kind of electronic support for the maintenance of medical records. The fear of automation expressed in the late 70s appears to have been replaced in the late 80s by a concern with the increasing

amount of documentation required in the practice of medicine. Most of the physicians with whom we spoke were eager to acquire computer support for their clinical practice because they believed that computers could help them to reduce the costs of paperwork, in terms both of time and of personnel, although they differed in the sophistication of their vision (see below). They saw computer assistance coming in two forms:

- Most believed an electronic medical record would be a great aid to managing patient data, although some were skeptical that paper records could ever be completely replaced, citing reasons of security, adequate access, and legal necessity.
- All felt that automated and timely capture of clinical data—the automatic transfer of data from laboratory computers and radiology reporting systems directly into the patient's chart without the need for manual transcription—would relieve them of much of the time they currently spend transcribing data and tracking down missing results.

The physicians' preoccupation with the data-management problem may have contributed to their lack of awareness of the way computers could improve their access to sources of medical information. While most acknowledged that online access to medical information and computer support of clinical decision-making could be useful, the perceived utility of such features was overshadowed by the crying need for a good clinical data-management system.

2.2. Fundamental Divergences

One of the major differences we observed in the clinics we surveyed was the degree of autonomy they enjoyed. The clinics we visited had varying degrees of control over factors in their environments that could have an impact on computer use. These factors included:

- *The physical plant:* Some clinics leased or owned the space in which they resided, and they therefore had control over its modification. In preparing for the introduction of an oncologist's workstation, these clinics would be able to rearrange their interior space, add more wiring, install additional telephone lines, and perform other modifications that might be required for workstation installation. Other clinics were subsumed by larger organizations and had little or no control over their physical space. These clinics would find it difficult to make room for new equipment or

to enhance other aspects of their physical plant in order to accommodate an extensive computer system.

- *Pathways to laboratory data:* Some clinics could choose the laboratories they employed to process tests; others were required to use the laboratory facilities of a particular institution. Clinics of the former type would be more likely to be able to negotiate with providers of laboratory services to establish electronic access to laboratory data than the latter type, which are at the mercy of a single service provider.
- *Patient-record storage:* Some clinics maintained their own patient-record rooms, while others used the patient charts maintained in the hospital with which they were affiliated. Clearly, clinics with control over the patient records would be able to transfer them to a computer-based medium of storage; other clinics would have to obtain institution-wide support or approval of a change in the format of the medical record.
- *Budget:* The private clinics we surveyed could allocate their expenditures as they saw fit, though the size of their budgets varied. Private clinics would therefore be able to budget for the purchase of new computer technology, while those clinics affiliated with universities or health maintenance organizations, having little or no control over their own budgets, would have to obtain the approval of their governing institutions.

We observed that the degree of control varied according to size, age, and organizational affiliation of the clinics. Private clinics, for example, had greater control over these factors than clinics affiliated with universities or health-maintenance organizations, and small, young clinics seemed better able to make major changes in their practices than larger, more established ones. However, while smaller clinics could be more responsive to change, they had fewer resources to allocate to new solutions and had a smaller base over which to amortize fixed costs of such changes. Although it would be appealing to argue that the incorporation of integrated workstations should facilitate the assignment of more resources and control to individual clinics and practitioners, especially in larger organizations, such changes are not likely until after the value of such technology has been well demonstrated at initial test sites. Questions of cost-effectiveness and legal liability or protection will be particularly pertinent in this regard.

2.3. Principal Barriers

The principal barriers to computer use that we noted in our survey had to do with the attitudes of users and with the nature of existing patterns of work. The wide diversity of computer sophistication among physicians and other clinical personnel has implications for the dissemination of the integrated oncologist's workstation. Great care must be taken to educate the entire user community (physicians, administrators, clerical staff) in the possible uses and the limitations of the systems that are installed. Sophisticated users will be able to make sophisticated use of the machine; naive users will be able to conceive of and make only limited use of the system (at least initially), and will therefore be less likely to embrace its use. Some practitioners who believe in the value of the technology in their offices will still rely on others to handle all computer interactions. While we acknowledge that this approach can provide useful benefits, many of the advantages of the workstation will best be realized through direct interaction by the care-giver or by the person who needs information stored in diverse databases. The experience of trained intermediaries to do bibliographic searching in medical libraries has demonstrated both the value and limitations in the use of surrogates for computer-based interactions.

As discussed earlier, potential sites vary in the flexibility with which they can absorb new technology. This variability among sites will force a number of tradeoffs in the design of an integrated oncologist's workstation and in the phasing of its introduction. While it would be desirable from a design and operations point of view to require uniformity across all sites, such a procrustean approach would risk providing suboptimal services for individual sites whose particular situations do not lend themselves to the standard solution. In particular, automated retrieval of laboratory data should be a fundamental component of the integrated oncologist's workstation, even though control of the pathways to laboratory and patient data will vary considerably. Furthermore, most potential workstation sites, especially large, well-established clinics or research institutions, will already have a pattern of work in place: how different tasks and responsibilities are assigned to different roles within the organization (e.g., physicians, nurses, administrators, data managers), how individuals in different roles communicate with one another, and how information is exchanged (e.g., who has access to different types of information and what docu-

ments are used to record and disseminate each type). Although the integrated oncologist's workstation will introduce many features that will change and improve the way people work, its impact must not be a brashly revolutionary one. If potential users perceive that the workstation will uproot existing patterns of work or modify them in a disruptive manner, they will not accept it.

The implication for the overall design of an integrated oncologist's workstation is that it must be as flexible and modular as possible. Compromises on these goals may simplify development and support of software but seriously jeopardize the acceptability of the tools to the intended users. The workstation should be designed as a toolbox of functions that can be assembled and integrated in many flexible ways to create systems that are finely tuned to the needs of individual sites, but which retain the uniformity of fundamental construction crucial to effective maintenance and growth. Although such an architecture can be complex in the development phases, modularity using well-defined standards for communication and interchange among system components can greatly simplify maintenance and the introduction of new features as the technology evolves.

C. A Brief Survey of Modern Computer Technology

There is a large and ever-growing array of computing technology that developers can use today to meet the needs of oncologists and others engaged in oncology-related work. As we discuss the design, functionality, and implementation of an integrated oncology workstation, we will be describing a system based on technology that is available now, though perhaps not yet in widespread use. Before we begin this discussion, it may be helpful to review briefly some of the trends in modern computing technology and to familiarize readers with some of the terminology we will be using. A more technical summary of the state of the art can be found in Appendices A and B.

At the outset, it is appropriate to explain our use of the term *workstation*. Throughout this document, we will be calling the system we describe *the integrated oncology workstation*, or simply *the workstation*. The use of this terminology may be confusing to readers who are familiar with the nomenclature used to describe different kinds of computer hardware: *microprocessors*, *minicomputers*, *mainframe computers*, *personal computers*, and *workstations*. These terms are used, often somewhat haphazardly, to classify computer hardware according to processor speed, memory size, the ability to support multiple simultaneous users, network capacity, and intended use. Our use of the term *workstation* is based less on any specific notions of computer hardware and more on the concept of a locus in which one's work is done. A pharmacist in his laboratory, sitting at a bench and surrounded by bottles of reagents, is at his workstation. Similarly, an oncologist sitting in front of a computer screen, on which is displayed the medical record of the patient in the examining room next door, the document describing the protocol on which that patient is enrolled, a page of a pertinent medical reference book, and the forms or other media through which medications can be prescribed and laboratory tests ordered, is at his or her workstation. Furthermore, because the computer screen is just a window into the functionality provided by one or more computers which may physically reside anywhere and which may be used from many locations, the oncologist's

workstation is not simply a physical place; it is a collection of tools and information that allows its users, be they physicians, nurses, data managers, clerical workers, or researchers, to do their oncology-related work.

The implementation and dissemination of this collection of tools and information will be greatly aided by the ever-increasing power of computer technology. The "better, smaller, and cheaper" trend continues in the computer industry, as more and more powerful machines become affordable for a wider range of potential users, including clinics and individual physicians. Processors continue to become faster, memory becomes denser, and mass storage continues to burgeon. Rotating magnetic media, such as *floppy disks*, *microdiskettes*, and *hard drives*, continue to grow in their capacity to store data and in the speed with which they permit those data to be retrieved. In particular, the enormous storage capacity afforded by *optical disc* technology (the same technology that has revolutionized the recording industry with the compact disc) is affecting the kinds of information being made available to users and the means by which this information is disseminated. For example, whereas earlier, more limited storage media allowed only excerpts of large documents to be made available on computers, optical discs allow publishers to store entire textbooks, complete with figures, on a single 4.5-inch diskette that can be read using off-the-shelf hardware.

The growing power of desktop machines in the 1980s has permitted the introduction of more powerful *operating-systems*, the low-level software that controls the machine and enables higher-level software to run. Although *multitasking* (the capacity to run more than a single program at once) and *virtual memory* (a technique that allows a computer system to simulate the large amounts of random-access memory [called *RAM*] that large computer programs often need in order to run) have been available on mainframe-, minicomputer-, and workstation-class machines for some time, these capabilities are now appearing in operating systems for personal computers.

Along with the growth in operating-system power has come an increase in the speed and flexibility of intercomputer communication networks. Local-area networks, wide-area networks, and online services have developed rapidly, and large-scale efforts are already under way to create the next generation of national research networks. Networking

technology has allowed machines of varying sizes and manufacturers to be linked to one another for purposes of remote access, data sharing, and cooperative computing. This technology has, for example, allowed hospitals and large clinics to develop departmental systems that are managed locally but which are still accessible to other machines and to disparate users throughout the institution. The national research networks, notably those of the Department of Defense and the National Science Foundation, allow users and computers around the country to work together using high-speed communications lines. Although only a small part of the health-care community (principally individuals at academic medical centers with technologically-oriented affiliated universities) has yet made use of these national networks to any large extent, the technology is clearly appropriate for clinical uses ranging from information dissemination to claims-form submission to the construction of centralized research databases.

The capacity to store enormous quantities of information of different types (e.g., text, images, and sound) and to retrieve it quickly has spawned new software technologies to support *databases*. In particular, the mathematical concept of a *relation* among various kinds of data has been used to support a highly flexible kind of database, called a ***relational database***, that makes it easy to retrieve data in many different ways, based on many different criteria. The relational database is rapidly becoming a standard paradigm for the storage of large quantities of data. Object-oriented database representation techniques exist for storing large objects such as text documents or images.

While relational databases permit users to establish relations between different *classes* of *data*, *hypertext systems* are exploring ways to permit users to establish relations among different specific *instances* of *information*. Hypertext systems, just now beginning to emerge from research laboratories, allow designers (and sometimes users) to make arbitrary *links* between arbitrary units of information (e.g., individual words, paragraphs, chapters) both within a single document or between different documents, much as footnotes and bibliographies allow authors to cross-link their scholarly writing and to tie their own work to the larger body of research literature. We strongly expect that many of the concepts derived from hypertext research will become an integral part of computer systems in

the 1990s. This expectation suggests some fascinating possibilities for new dynamic and cross-linked forms of information bases such as those currently provided by PDQ and CANCERLIT.

Modern computer technology has introduced many new ways for humans to interact with computers. On the output side, an ever-increasing array of hardware and software tools is supporting the use of high-density bitmapped graphics, windowing environments, high-resolution displays, and laser printers to present information to users. On the input side, software such as on-screen menus, icons, and buttons, and hardware such as pointing devices, touch screens, and even speech-recognizers are being employed to provide users with new ways of communicating commands and information to the computer. Software architects are employing these technologies to create user interfaces that are easier to use and that give users access to greater functionality than line-oriented, typed-command interfaces. Our experience has shown these advances to be of particular importance for systems designed for use by physicians, because physicians are often strongly resistant to typing. Designers have begun to explore new ways of providing users with fast and flexible ways of navigating through large amounts of information or among multiple, simultaneous applications. The well-known desktop metaphor, for example, assists users in making sophisticated use of computers quickly and with little training. Other interface designs have begun to exploit the multimedia possibilities afforded by high-density color screens, gray-level images, high-quality sound reproduction, and the integration of computers with video-display technology such as videocassette and laser disc players, and 3-dimensional graphics.

In this environment of increasingly powerful hardware and software, designers exploring the application of techniques derived from artificial intelligence research to problems in the domain of medicine have been able to provide clinical users with powerful systems that could once run only in research laboratories. Some of these clinical advice systems can aid physicians in making diagnoses; others can help them to create therapy plans that follow the treatment guidelines specified in complex clinical trial protocols. The amount of knowledge these programs need to perform their tasks is very large, and the strategies, or *algorithms*, they use are complex. Nevertheless, some of these programs are becoming suffi-

ciently mature that the 1990s will find clinicians using them on an increasingly routine basis.

It is with this ever-growing array of hardware, software, and networking technology that we expect the integrated oncologist's workstation to be built. In our discussion of the workstation we explicitly avoid reference to specific hardware and software. Our restraint is based on our own considerable experience in designing and building state-of-the-art systems. We have learned that computer hardware is highly transitory; it is guaranteed to change rapidly as new technological breakthroughs are achieved and brought to market by different vendors. Likewise, computer software, such as operating systems, languages, database systems, interface technologies, and applications have undergone significant improvements, sometimes driven by and sometimes driving the evolution of computer hardware. The policy within our own research group has been to avoid designing and building advanced systems for today's hardware; by the time a project is finished, the hardware and software for which it was targeted would very likely be obsolete. Instead, we have tried to anticipate the trends in basic technology that are likely to evolve over time. We have learned to avoid committing to specific vendors, as individual products and companies may have remarkably short lifespans in the high-technology marketplace. We have tried instead to conform to evolving industry- and community-wide standards, not to the *de facto* standards that arise because of the temporary dominance of a particular vendor's products, but to standards that are built on a principled examination of needs and that have a national or international scope of acceptance. The most expensive part of developing a modern computer-based product is no longer the cost of the hardware, but rather the cost in human effort required to develop it and to keep it running in changing hardware and software environments. If the design of an integrated oncologist's workstation is based on standards that support its essential functions and human interface design, it will be insulated from transient hardware and software and will therefore be more likely to be durable.

D. Scenarios of Clinical and Research Use

Before beginning a discussion of the specific features and components of the integrated oncology workstation, we believe it is useful to set the scene by describing some of the environments in which the oncologist's workstation might be used. Accordingly, we have prepared three imaginary scenarios that show the oncology workstation being used by oncologists in a small private clinic, by a data manager for a fictitious research group, and by the research director of that group. The scenarios and the figures that accompany them will help to provide a context for the detailed discussion in section II.E.

A Small Private Clinic

Dr. Enderby and Dr. Hettrich have a private oncology practice in a small but prosperous town in northern California. Their clinic occupies part of the second floor of an old hospital and includes a waiting area for patients, a reception/administration area, two examination and treatment rooms, a small laboratory room, and an office for each physician. Drs. Enderby and Hettrich employ a chemotherapy nurse, a laboratory technician, and two clerical staff. The physicians acquired the integrated oncology workstation a few years after starting their practice. They have 4 graphical terminals, one at the main reception desk, one at the nurse's station, and one in each physician's office. All are connected to a workstation in the clinic's laboratory room.

The practice has its own analyzer for doing routine blood tests, and this machine is connected to the workstation. Test data are transmitted directly to the workstation for inclusion in patient records. The workstation provides a flexible interface to many kinds of laboratory equipment using standard data-transfer protocols. The blood analysis machine could not be connected initially, but the workstation's interface module was augmented to include it with little difficulty once the workstation was installed.

On Friday, March 16, Dr. Jones' office calls Dr. Enderby's office to arrange an appointment for Sally Wong, who has been recently diagnosed with colon cancer. Dr. Enderby's nurse, Joanne Young, takes down the

demographic information, including the fact that Ms. Wong had recently been an inpatient at Local Care Hospital. Nurse Young establishes a computerized patient record for Sally Wong and requests that the workstation contact the hospital computer to retrieve any relevant clinical information.

To do this, she selects the "Patient Records" button on her desk, and opens a window that contains a menu of options. These options include retrieving active patient records and archiving inactive ones, printing patient records, and creating new patient records. She selects the option for creating new patient records, and is presented with a form. On this form is a space for recording other institutions that are keeping electronic records for this patient. When she selects the "Choose Data Sources" button, Nurse Young is presented with a form that contains a menu of institutions to which the Enderby and Hettrich clinic has electronic access. On this form, she can select the name of an institution from a menu and then enter the number by which the patient is identified at that institution (Figure 1). Nurse Young selects Local Care Hospital from the menu and enters Sally Wong's hospital identification number. Nurse Young also knows that Dr. Jones frequently uses Alpha Laboratories for laboratory work, so she instructs the workstation to contact Alpha Laboratories as well for possible information on Sally Wong.

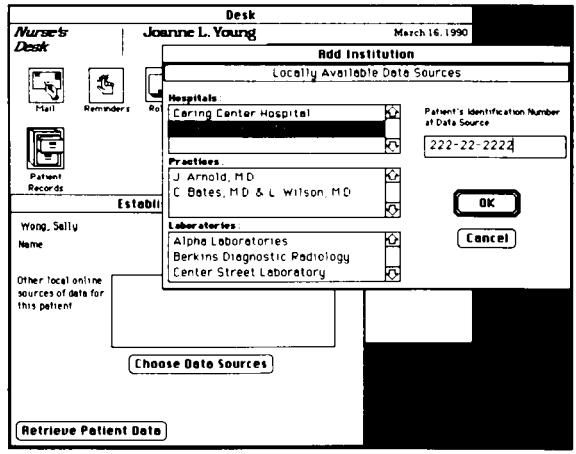

Figure 1:
The nurse in Dr. Enderby's office indicates that data regarding Sally Wong are available from the central computer at Local Care Hospital.

Nurse Young then selects the "Retrieve Patient Data" button. The workstation provides feedback as it connects to the Local Care Hospital system, via the network facilities Drs. Enderby and Hettrich have recently added to their workstation. To guarantee secure access to the hospital's medical records, the hospital's data server asks Nurse Young for the hospital password assigned to the Enderby and Hettrich clinic. The hospital computer checks first for Dr. Enderby's clearance to the hospital records, then for his clearance to Sally Wong's record. When Dr. Enderby's access clearance has been established, the hospital computer permits the oncologist's workstation to present Nurse Young with a list of the classes of data it finds in Sally Wong's record (surgical reports, pathology reports, nursing notes, and others) and asks Nurse Young to select the kinds of data she wants to retrieve. She selects surgical reports, pathology reports, laboratory reports, hospitalization reports, and radiology reports, in accordance with standard practice at the Enderby and Hettrich clinic. When she has finished, the workstation informs her as it retrieves the data (Figure 2) and notifies her when it is through. The workstation follows a similar procedure as it retrieves data from Alpha Laboratories; although there is less data to retrieve, the process is somewhat slower, because Alpha Laboratories is accessible only via modem.

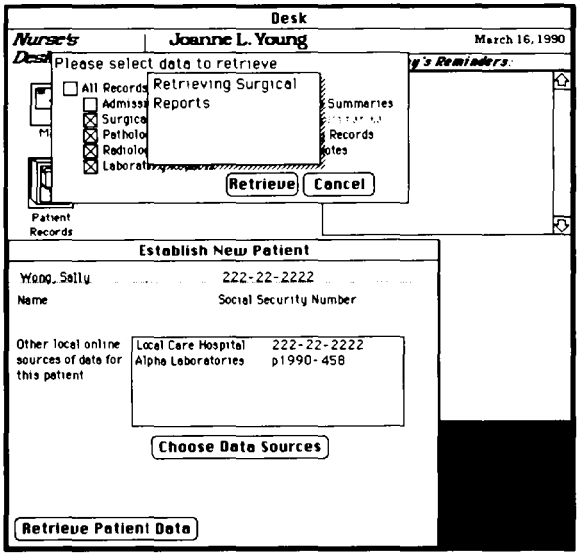

Figure 2:
The workstation keeps Nurse Young informed on its status as it retrieves Sally Wong's data from Local Care Hospital.

Patient Chart Browsing

Because Nurse Young retrieved Sally Wong's medical records on Friday, they are available on Monday morning, when Dr. Enderby comes into his office, looks down his schedule of patients on the workstation, and sees that Sally Wong will be seen as a new patient. Dr. Enderby knows from his conversation with Dr. Jones last week that Ms. Wong is coming for consideration of adjuvant chemotherapy for Duke's C carcinoma of the colon. Dr. Enderby sits down at his oncology workstation to review the information presently available about Ms. Wong (Figure 3).

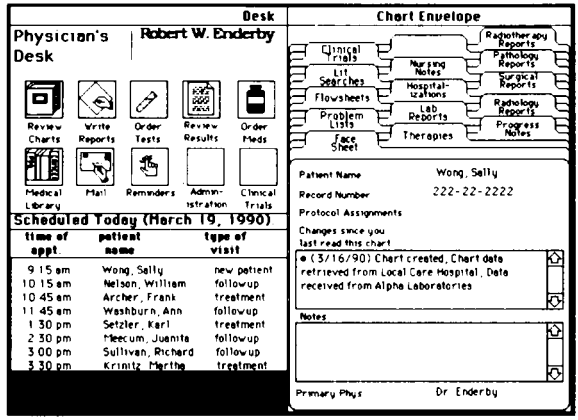

Figure 3:
Using the mouse pointing device, Dr. Enderby selects Sally Wong's name from his day's schedule and is shown a graphical representation of her recently created chart. He opens record sections of interest by clicking on the pertinent chart "tab" with the mouse.

For general medical background, Dr. Enderby first looks at the admission history and physical from her recent hospitalization and at her discharge summary. Dr. Enderby sees that in general Ms. Wong had been in excellent health but that recently carcinoma of the colon had been diagnosed. He reviews her hospital record to obtain additional information regarding the stage of her disease. First he reviews her operative report (Figure 4) and her radiology reports, which include a chest x-ray and a CT scan of abdomen and pelvis. He reviews general chemistry and CBC results, and then goes to the surgical reports to read the operative report for surgical staging. He finds that at the time of laparotomy all the apparent disease was confined to the bowel. He then opens the surgical pathology reports to review the pathology from the patient's laparotomy. He finds that a partial colectomy was performed and that a 2.5 x 3 cm mass was found in the area of the sigmoid colon. He also notes that there

was a 5 cm tumor-free margin of normal colon on either side of the lesion. The diagnosis was moderately differentiated adenocarcinoma, and extension into the pericolic fat was noted. He also notes that 2 out of 16 mesenteric lymph nodes were found to contain metastatic carcinoma.

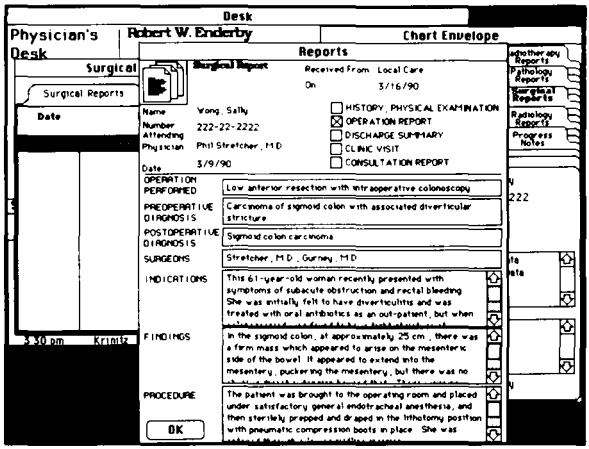

Figure 4:

Sally Wong's surgical report is displayed when Dr. Enderby selects the corresponding tab on the graphical patient record.

Literature Searching and Protocol Eligibility

The information in Sally Wong's medical records convinces Dr. Enderby that the patient has stage C adenocarcinoma of the colon, which was completely resected during recent surgery. Selecting the "Clinical Trials" tab from Sally Wong's chart, he goes to the workstation's protocol-eligibility scan to see if the patient is eligible for any of the cancer clinical trials active within the oncology groups with which he and Dr. Hettrich are affiliated. The workstation analyzes the pertinent patient data from its database and scans the eligibility tables of the available clinical trials. After identifying appropriate clinical trials, the workstation presents the results of the scan to Dr. Enderby.

Dr. Enderby notes that among the protocols for which Sally Wong is presently eligible is an NCI high-priority clinical trial. He selects the name of this protocol from the list of scan results, and the workstation presents him with the information about this trial contained in the PDQ database, which resides on a CD-ROM connected to Dr. Enderby's machine (Figure 5). After browsing through this information, Dr. Enderby asks the workstation to print out the patient information statement for colon cancer from the PDQ database for Sally Wong.

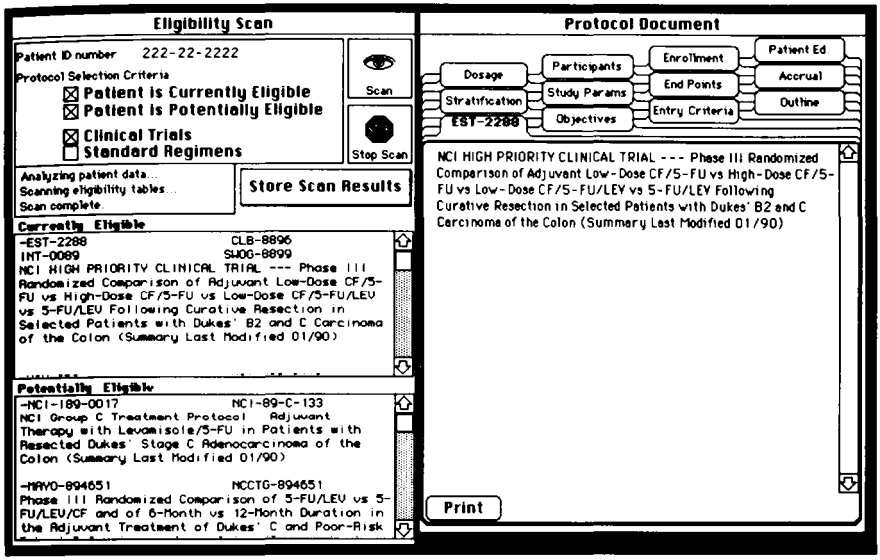

Figure 5:

The workstation has identified a high-priority NCI trial for which the patient is eligible. When Dr. Enderby requests more information by clicking on the protocol identification number, a full PDQ-style description of the protocol is loaded from CD-ROM and displayed as a graphical document for perusal.

Protocol Enrollment and Treatment

Dr. Enderby then goes to see the patient, performs a complete history and physical, and discusses clinical enrollment with her. Dr. Jones had discussed the possibility of the need for chemotherapy with Ms. Wong, and she has come to the clinic prepared to give her consent to protocol participation. After Ms. Wong has read the PDQ patient information statement Dr. Enderby gives her, she agrees to be enrolled on the high-priority intergroup adjuvant colon cancer protocol.

Having had Ms. Wong sign the human subjects consent form, Dr. Enderby returns to his oncologist's workstation in order to electronically register and randomize her through the cooperative group protocol office. Dr. Enderby selects the enrollment section of the protocol from the protocol text browser (Figure 5) and notes that online enrollment for this protocol is available. He therefore instructs the computer to register the patient and obtain a randomized assignment. After a few seconds, the computer returns with information that the patient has been registered and randomized to arm IV of the protocol (Figure 6).

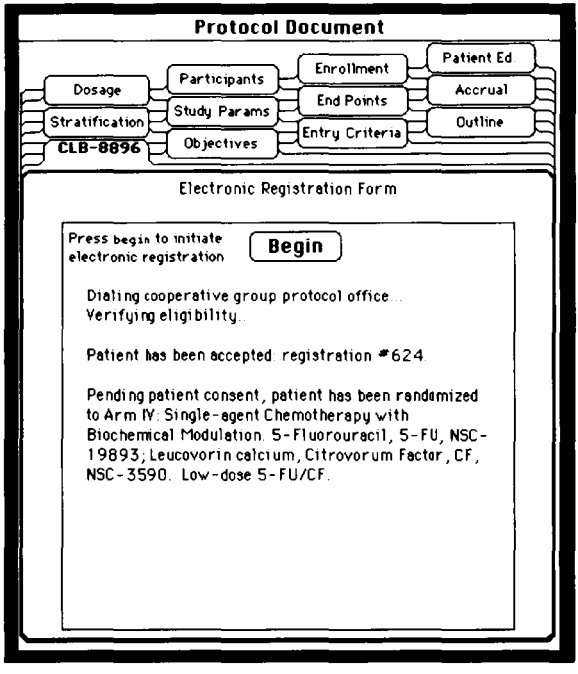

Figure 6:
The workstation automatically dials the central study computer in order to enroll Sally Wong in the clinical trial and to report to which arm she has been randomized.

By the time the patient is registered and randomized, the CBC obtained at Dr. Enderby's office that morning is available. Dr. Enderby selects "Therapies" from the patient's chart and, by selecting the "Recommended Protocol Therapy" button, instructs the workstation to calculate starting doses for the first cycle of protocol therapy. When proposed doses are displayed (Figure 7), Dr. Enderby examines them and agrees with the recommendation. When he selects "Record Prescription", the dosing information is relayed to his clinic pharmacy. The computer system recognizes that there are investigational drugs involved in this protocol and that only one bottle of investigational drug is available. The computer asks Dr. Enderby if he wishes to order more drugs, and when Dr. Enderby indicates yes, the oncologist's workstation automatically generates an electronic mail message to the National Cancer Institute, Cancer Therapy Evaluation Program, Investigational Drug Branch, to order additional drugs for this protocol.

Test Ordering

The advice program examines the documentation for this protocol, and reminds the doctor to order several laboratory tests prior to the next

protocol treatment (see "Note" in Figure 7). Dr. Enderby orders the tests by selecting the "Order" button on the workstation's recommendation screen. The workstation knows in which laboratory each test is usually performed, so it prints an order slip stamped with a bar code for the complete blood count on the printer in the clinic's laboratory room. The laboratory technician will use the bar code printed on the order slip to configure the local equipment to perform the required CBC and to report the results to the workstation with Ms. Wong's medical record number.

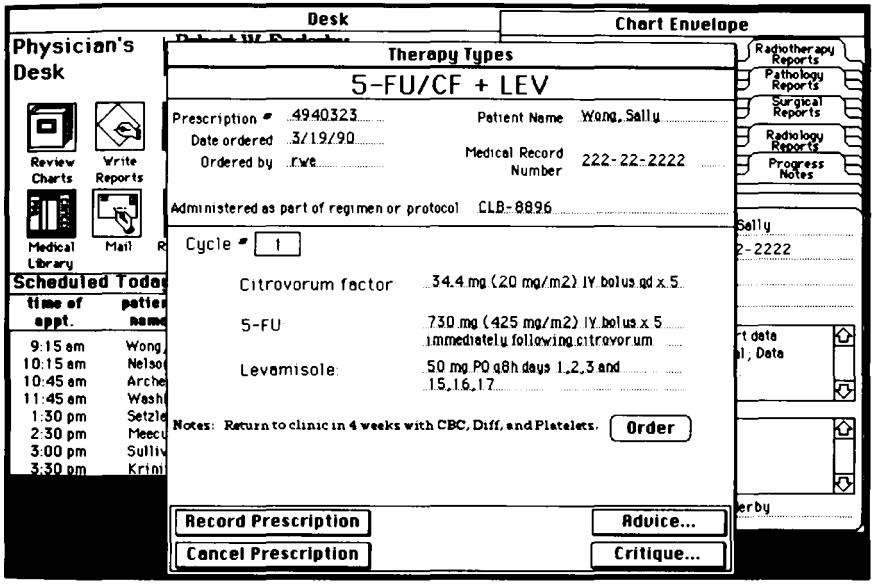

Figure 7:

The workstation displays dosing information appropriate for Sally Wong's therapy plan, blood counts, and body surface area.

The workstation sends an electronic order to Alpha Laboratories for some of the other tests that Dr. Enderby chooses to order for Ms. Wong. There are two private laboratories in town, Alpha and Biologica. Drs. Enderby and Hettrich used to send most of their laboratory work to Biologica, but when they installed the workstation, they switched to Alpha because it offered direct computer access to laboratory results. The workstation in the clinic is set up to dial a computer at Alpha Laboratories twice per day and to retrieve any laboratory-test results that are pending.

The program then sorts these results, and inserts them in the appropriate patients' medical records. Drs. Enderby and Hettrich can use the workstation to print daily reports of the laboratory results that have been received, and the workstation also notifies them of new laboratory data when they first open a patient's chart for review. Within a year, Alpha Laboratories plans to offer automated transmission of laboratory results; then data will be transmitted directly to the clinic workstation as soon as the tests are performed. Drs. Enderby and Hettrich still send some work to Biologica, which has its staff call in results over the telephone or sends them via a courier service. These results are then transcribed into the workstation by one of the clinic clerks.

After Ms. Wong has received her treatment and is checking out of the clinic, the receptionist uses the workstation's scheduling program to book her next treatment visit. At the end of the day, the workstation will export information about Ms. Wong's visit to the clinic's administration and billing system for use in preparing bills, insurance forms, and other administrative documents.

Critiquing

While Dr. Enderby is seeing Ms. Wong, Dr. Hettrich is treating another patient, Paula Fenton. Ms. Fenton has Hodgkin's disease but is responding well to BAOG protocol 8H-85-2, under which she receives alternating courses of MOP(P) and ABVD. Today Ms. Fenton is scheduled to receive cycle 2A of MOP. After reviewing the patient's flowsheet, Dr. Hettrich uses the workstation to order MOP, prescribing 6 mg of nitrogen mustard IV day 1, 2.0 mg of vincristine IV day 1, and 150 mg of procarbazine PO days 2-15.

Dr. Hettrich often prefers to prescribe treatments herself, and then to ask the workstation to critique her treatment plans. When she has finished creating a MOP(P) 2A treatment plan, therefore, she selects the "Critique" button on the prescription form. The workstation soon reports that the 6 mg dose of nitrogen mustard may be too low (Figure 8). The patient's 100% dose is 8 mg, and, given her blood counts and toxicity profile, no dose attenuation is warranted by the protocol. The workstation displays a reference to the protocol document; by selecting this reference pointer, Dr. Hettrich is presented with the appropriate page in the protocol document. She reviews the document and, finding she agrees with the critique, changes the doses and records the prescription.

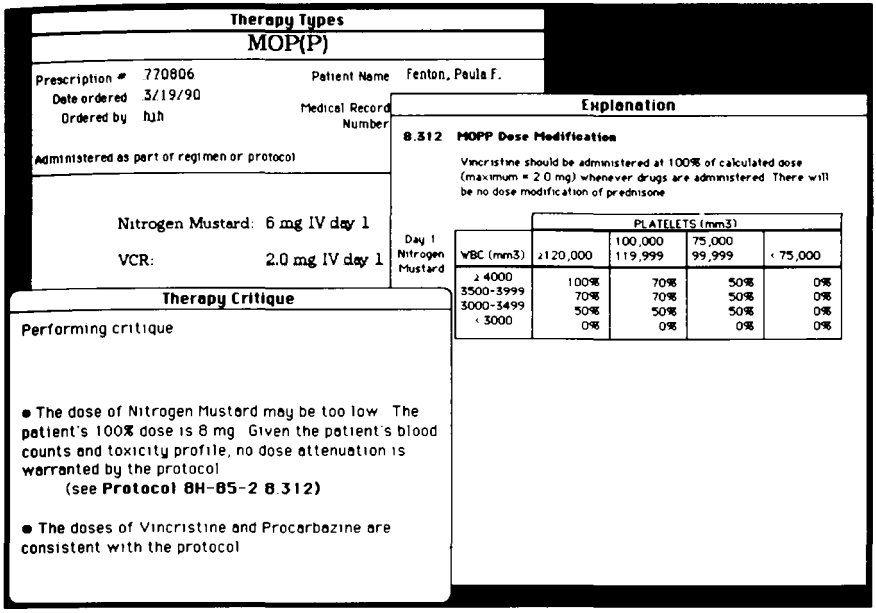

Figure 8:

The workstation provides a critique of Dr. Hettrich's proposed treatment plan. It points out that the protocol would normally call for a higher dose of nitrogen mustard and, on request, displays the pertinent page from the protocol document.

A Data Manager

Ms. Kirk is a data manager with BAOG, the Bay Area Oncology Group. BAOG sponsors approximately twenty clinical trials in northern California, and Ms. Kirk is responsible for managing the data being collected for seven of these trials at a dozen clinics. The arrival of the integrated oncologist's workstation has altered Ms. Kirk's work patterns considerably. When workstations began to appear in oncologist's offices, Ms. Kirk immediately noticed an improvement in the organization of patient records—it became much easier to track patient data in the well-organized, laser-printed charts the workstation produced. When BAOG began to distribute its protocols to workstation users on CD-ROM a few years ago, first as browsable text and then as knowledge bases for the workstation's decision-support program for protocol management, Ms. Kirk noticed that the quality and completeness of the data began to improve at workstation sites.

At that time, Ms. Kirk learned to use the protocol-management features of the workstation herself, which made it much easier for her to examine patient data for missing or suspicious data values. Using these tools, Ms. Kirk can extract data from patient records and create her own databases for each protocol (Figure 9). For offices without workstations, this still means going to the clinic, transcribing the data from patient charts by hand, and then returning to her office to copy the data into her workstation. At some workstation sites without adequate modem connections, Ms. Kirk must still go to the clinic and copy onto diskettes the data for the patients she is tracking (or have the clinic mail diskettes to her), but most workstation sites have modems, and she is able to log on to them from the workstation in her office and download encrypted patient data directly. The workstation's tools help her scan the data for missing values and for values that are unusual or out of range. Ms. Kirk is also able to forward complete patient records or composite summaries electronically to the study chair for review, and to the NCI for periodic interim reporting of study outcome.

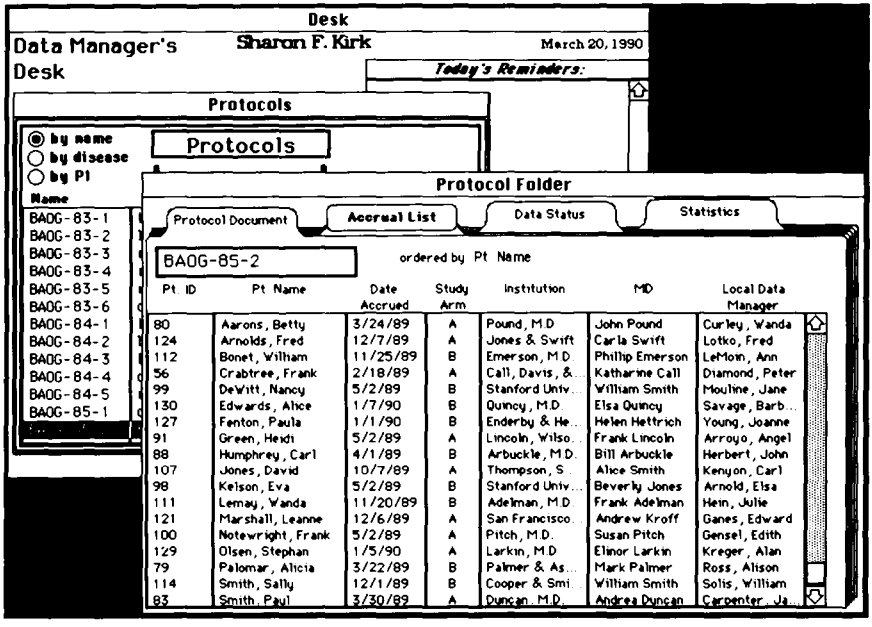

Figure 9:

The BAOG data manager examines the accrual list for one of the protocols that she oversees. The workstation supports her data-accumulation and analysis tasks through a series of software options that she selects with the mouse.

An Oncology Group Research Director

Dr. Campbell is the chair of BAOG, where one of his chief concerns is the promulgation of BAOG's clinical trials. He is always looking for ways to increase enrollments in their active protocols, both by making prospective patients and physicians more aware of the trials they sponsor and by making it easier for oncology clinics to participate in them. Furthermore, the rising costs of clinical trials research has made it increasingly important to create protocols that have well-focused research goals, so Dr. Campbell has been investigating ways to ease the design of clinical trials.

Several years ago, BAOG installed the integrated oncologist's workstation at their headquarters, using a configuration that emphasized the protocol data-management features of the clinical workstation and that also included tools for the design and creation of new clinical trials. One of their first uses of the workstation was to transcribe the full text of all of their active protocols into the machine, using the text-organization and hypertext tools of the workstation to make the protocols easy to search and browse. With the text stored in the workstation this way, users can do fast searches over the full text of any or all the protocols BAOG sponsors, as well as browse easily through the protocol descriptions, using the hypertext links to skip quickly among chapters and sections and to follow cross-references. Currently, BAOG publishes its protocols on a CD-ROM that contains protocols for ten other oncology groups in the western and southwestern United States. Publication is handled through a cooperative publishing agreement sponsored by a major publishing house in Los Angeles.

At the same time BAOG was transcribing their existing clinical trials into hypertext format, Dr. Campbell and his research colleagues were using the workstation to augment the textual descriptions of the protocols with knowlege-base components and rules for eligibility. The workstation provides tools for easily encoding the treatment guidelines and the therapy and test-ordering schedules into charts, tables, and rules (Figure 10). These structures can then be used by the workstation's clinical advice components to help clinicians manage the administration of the protocol by monitoring treatment schedules and doses and by assisting them in ordering laboratory tests required by the clinical trial. The protocol development tools on the workstation make it possible to link the knowl-

edge base structures used by the advice components to the textual descriptions of the protocols; thus clinical users may relate the suggestions generated by the advice components to the text of the corresponding protocols. BAOG is working on creating knowledge bases for all their active clinical trials.

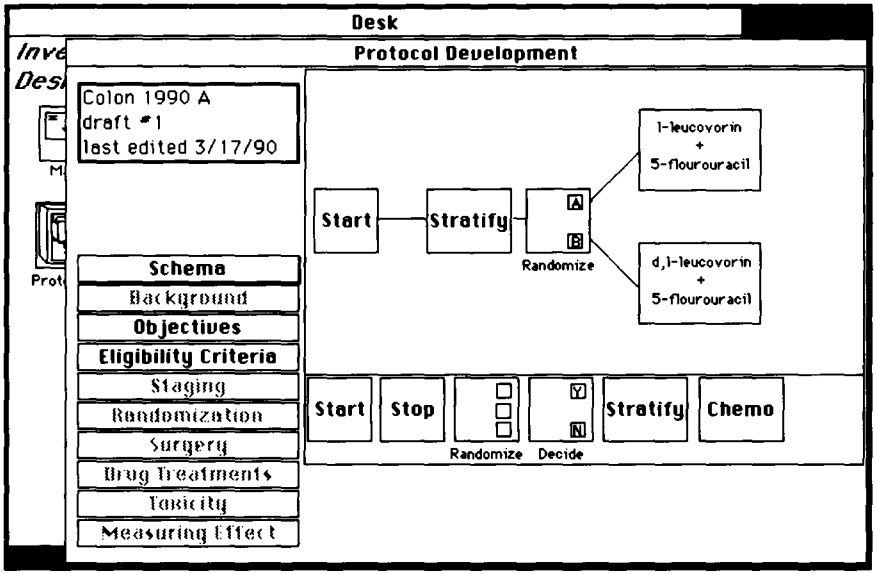

Figure 10:

Dr. Campbell is using the workstation's protocol-development tools to encode a newly developed protocol for colon cancer. The workstation automatically converts the drawing he produces into programs that can later be used for decision support during clinical trials.

The tools to generate protocol knowledge bases are integrated with other workstation tools to develop databases for managing and tracking patient data for the clinical trial. For example, the tools Dr. Campbell and his colleagues use to create a knowledge base of treatment guidelines for a protocol also allow them to specify the parameters they wish to collect in the clinical trial. This list of parameters can subsequently be used by the data collection programs Ms. Kirk and other data managers employ to extract data of interest from patient charts. The workstation itself also contains many statistical tools that aid in the analysis of clinical trials data. However, although BAOG plans to phase out its old computer system

within several years and to switch entirely to a network of integrated workstations, BAOG still uses the old system to perform a number of analyses. The data gathered and tracked on the workstation must therefore be transferred to the BAOG mainframe. Because of the workstation's flexible networking capabilities and its reliance on industry standards for data representation and storage, exchanging data between the two systems is not difficult.

Some of the latest software that has become available for the integrated oncology workstation is a set of software tools to aid in the design of new clinical trials. These tools are an extension of the toolset for creating protocol knowledge-bases, and they help the protocol designer develop hypotheses, create statistical projections, estimate required sample sizes, design balanced research arms, and perform other tasks in the process of designing a new clinical trial. The toolset allows designers to design a new trial from scratch or to adapt an existing trial to investigate new variations.

E. Discussion of a Clinical Oncologist's Workstation

In this section, we describe the oncologist's workstation in detail. The description is nontechnical, concentrating on the organization of the workstation's components and the features each one should provide. The integration of all workstation components is a primary concern. The patient data entered in one module, for example, should be available to every other module, with the workstation maintaining a unified clinical database for all data. Furthermore, every module should be linked for easy access from every other module: A physician using the patient chart module to review a patient's record should be able to consult with PDQ or to bring up the protocol eligibility module to screen the patient for inclusion in a protocol without having to leave the patient's chart. This integration of data and functionality across the workstation has ramifications both for the design of the workstation's database and for the design of its user interface; workstation developers must accordingly bear in mind the need for this integration and compatibility as they design each component of the system.

E.1. The Master Screen

Systems that offer more than a single service must have a central coordinating point, an entry point into all the functionality that the system provides. Designers of integrated systems have given this coordination center many names: *master screen*, *home screen*, *entry screen*, and *main menu* are just a few. As an integrated system, the oncologist's workstation should provide such a jumping-off point to the different services it provides. Because the oncologist's workstation will offer different services to different kinds of users, its design should allow its master screen to be tailored to each class of user, and even to each individual user. The master screen for a physician should look considerably different from the one for a data manager, for example.

Although the master screen should serve as the entry point to all of the workstation's functionality, users should not be *required* to return to the master screen each time they wish to use a different module. The

oncologist's workstation should be designed so that modules can run simultaneously and can be activated easily from one another. In the discussions of individual components below, we will indicate how buttons, menus, and other navigational devices should be used to facilitate quick access to other modules.

E.2. Functions for Maintaining Electronic Medical Records

The availability and detailed features of these functions should vary according to the role of the user. Physicians, for example, might be able to prescribe medications and therapies but nurses might not, and administrators might be able to see the demographic data in a patient's chart but not the clinical data. The workstation design, however, should allow the capabilities assigned to each role to be customized for each installation since not all practices will wish to adhere to the same access conventions.

The remainder of this section is organized around the different tasks that a clinical user might want to perform with a patient's chart.

2.1. Functions for Browsing

One of the primary features of computers is their ability to store and retrieve data. Later in this section we describe the specialized tools an integrated oncology workstation should provide to allow physicians and others to store patient data in electronic form. Just as important as these tools, however, are those that permit users to retrieve stored patient records, to review them on the workstation's screen in different forms, to modify or add to them, and to print them in a variety of formats. Many computer-based record systems present information on the screen in much the same way as they print it on a printer—by displaying it in one continuous stream from beginning to end. Often the only control a user has over this stream is the ability to stop it and start it, and sometimes to "page" through the output by pressing keys on the keyboard that display a page of information at a time. This is not a very efficient way for people to examine information; it is analogous to a reader having to start at the beginning of a book and to turn every page, one at a time, to get to page 74, where the information the reader wants to see is located. Further-

more, if he were trying to compare the information in two books presented this way, the reader would have to close the first one before opening the second.

Modern computer systems have enhanced their use of the screen to provide users with a much more flexible way of looking at information, called *browsing*. When browsing through information, users are not constrained to look at it linearly, from beginning to end; they can skip around, looking at the information in any order that meets their needs, just as they would flip through the pages of a book or a medical chart (see the scenarios in Section II.D for illustrations of this concept). Furthermore, they can review more than one document simultaneously, just as they can when they have more than one book or more than one section of a chart open at a time. The tools of an integrated oncology workstation should be particularly designed to allow physicians and other health-care professionals to browse through patient charts and other sources of medical information.

2.1.1. Review a Chart

This general function lets users browse the data constituting a patient's medical record. Patient records usually consist of several classes of data:

- textual but formatted demographic data
- quantified results of laboratory tests
- quantified records of the doses of medications given
- textual notes and reports
- graphical reports

When patient data are stored in paper records, they are often recorded on different forms that support paper-based storage and retrieval of different classes of data: flowsheets for time-oriented, quantified laboratory and dosing data, and face sheets, problem lists, and progress notes for different kinds of textual data. Graphical reports such as EKGs or electrophoresis tracings are generally inserted verbatim in the chart with some summarization statement in the progress notes. Because the user interface should conform as closely as possible to the paper instru-

ments with which users are familiar,[1] the workstation should store patient records in a standard database whose schema permits the efficient retrieval of data for display in these formats.[2]

Users should be able to open a chart for browsing either by selecting a patient identifier from a menu or a list or by typing it into a selection field. Patient identifiers will typically be the patient's name and medical record number, but the workstation should permit them to be tailored to individual practice preferences. The workstation should be able to impose different orderings on the presentation of lists of patient identifiers so that users may find the patient they want more easily. Users might want to see patients listed in alphabetical order, for example, or by the date of their next visit to the clinic.

The specific components of a patient chart will vary from site to site, and the workstation design should be flexible enough to accommodate site-specific modifications to the chart's layout and contents. In general, however, the workstation interface will present the following chart sections:

- face sheet
- flowsheet
- problem list
- reports (progress notes, radiology, laboratory, and surgical reports, and others)

Each of these chart sections is described in more detail below.

---

[1] Many papers in the user-interface literature have discussed the value of user interfaces that mimic the "real world." People are able to apply their knowledge of real-world objects and the relationships among them to interactions with computer representations, so contemporary design philosophy suggests that whenever feasible or applicable, designers of user interfaces should incorporate real-world manners of presentation and interaction into the computer interface: text appears on simulated pages, numeric data are entered via calculator-like keypads, and so on. There may be times when simulation is not feasible, when it complicates the user's interaction with the machine, or when it complicates attempts to deliver a desirable functionality. The latter is especially true in situations in which the computer offers functionality that had not been anticipated or deemed feasible before. The designers of the workstation must therefore carefully weigh the costs and benefits of simulating real-world paper records.

[2] The patient database schema must also support the efficient retrieval of data by other programs that may make different kinds of queries. Decision-support systems, for example, may need to know more complex relationships among temporal data than strict sequence. Refer to the technical appendices for a discussion of these requirements.

In addition to displaying the contents of a patient's chart, the workstation should be able to provide users with a summary of the changes that have occurred in the chart since the last time they looked at it: the arrival of test results, the administration of medications, the reporting of new findings, and so on. The workstation interface should provide buttons or menus that allow users to move quickly among the chart sections, and it should permit as many of the sections to be viewed simultaneously as is possible.

1.1. Browsing Chart Sections

Browsing Face Sheets

The face sheet should contain demographic and administrative data about the patient. This section of the chart may be the only section accessible to non-clinical workstation users. The workstation should provide utilities that allow individual clinics to design face sheets and other forms to suit their own needs.

Browsing Flowsheets

The workstation should be able to display flowsheets in a variety of formats. Users should be able to specify the parameters they wish to see in the flowsheet, how those parameters should be organized (e.g., how they should be organized into sections and the order in which they appear), the number of columns of data to display, and the granularity of each column (e.g., whether each column of the flowsheet represents the time something was entered into the patient's record or some regular interval of time such as a day, a week, or a month). The workstation should allow users to save these flowsheet specifications and to use them to display data by selecting them from a menu. Clinics and individual users should thus be able to create libraries of flowsheets tailored to different purposes.

The workstation's flowsheet mechanism should allow users to see more than one section of the flowsheet at a time. When there are more data in the record than can fit on a single flowsheet page, the user should be able to shift the data in the flowsheet's columns left and right, like a person reading a scroll (this action is often called *scrolling).* If a user selects a cell within the flowsheet, the workstation should, if the cell is empty, present a menu or a window that allows the user to record a new

value. If the cell is not empty, the workstation should present a window that summarizes the initial entry of the value (see the section on reporting complaints, findings and diagnoses below), so the user can review the data value in context and, if he wishes, modify it.[3]

Browsing Problem Lists

A problem list is a way of documenting on paper the relationships among the complaints, findings, diagnoses, and treatment interventions that are recorded in a patient's chart. Because the workstation's functions for reporting clinical information (described below) allow users to establish relationships among different pieces of clinical information as they are reported, the workstation could derive a problem list automatically.[4] To see a problem list, users should be able to select the kinds of clinical information they would like to see in the list (chief complaints, findings, diagnoses, toxicities, interventions, etc.) and the kinds of links they would like to see among them (*causes*, *relieves*, etc.).[5] The workstation should use these instructions to display the requested information on the screen (or to a printer) in a format that clearly shows the relationships among the data. Individual users or workstation sites should be able to save particular sets of specifications (e.g., "show all chief complaints, diagnoses, and hospitalizations chronologically, and show their relationships to one another and to treatment interventions") as *problem list templates*, so that users can quickly review problem lists in different formats, depending upon particular needs.

Browsing Reports

The workstation should be able to display a list of all of the reports and progress notes in a patient's record, sorted by type, date, or both, and users should be able to examine any report by selecting it from the list. As

---

[3] Modifications to established records will of course require suitable retention of the original information and attribution for any changes. The electronic record and flowsheets must at least be no more prone to retrospective tampering than is the paper record of today. Properly designed, it can be even more resistant to such changes without proper documentation.

[4] Most care-givers have personalized approaches to the creation and use of problem lists. The discussion here is not meant to suggest that such personal styles should no longer be supported. Rather the computer might simply propose a problem list for editing and further enhancement by the system's user prior to approval.

[5] Readers familiar with databases will recognize this as specifying the parameters for a database query, whose results will be displayed as the content of the problem list.

described in the section on writing a note, below, clinics should be able to design their own formats in which to record and display reports, though the workstation should provide a library of standard formats. Computer-based storage of graphical reports such as EKGs can be handled digitally on conventional media, but this is not likely to be a practical approach for storing large amounts of graphical patient data. Newer optical disc technologies (for example, CD-ROM and videodiscs) provide practical approaches to the efficient and cost-effective storage of such data for rapid retrieval.

2.2. Functions for Performing Clinical Tasks

An integrated oncology workstation should help physicians document many of the clinical actions they perform. Physicians and other qualified medical personnel should be able to use the workstation to write orders for laboratory tests or other diagnostic procedures and to prescribe medications and other therapies.[6] The test results and treatments that users record for a patient should become part of the patient's record and should, therefore, be available to other modules. As we have just seen, the chart-browsing module should be able to display them; other modules, such as billing and insurance modules (whether on the workstation or a secondary financial machine), should also be able to extract information from the patient's chart for their own purposes. Laboratory data should be accessible for review by clinicians and for use by decision-support programs that generate patient-specific therapy advice or that screen individual patients for protocol eligibility.

When possible, however, the workstation should do more than simply document clinical actions: it should also help perform them. When a physician orders a laboratory test, the workstation should print the laboratory order forms that accompany samples, plus other printed matter. If the workstation is networked to laboratory computers that can receive electronic orders, the workstation should be able to transmit an electronic order for the test the physician has ordered to the laboratory that performs the test. A workstation in a particular clinic should maintain a table of information about all the tests and procedures that can be ordered at that

[6] As previously emphasized, the functions for performing clinical actions should be available only to the individuals for whom this degree of clinical responsibility is appropriate. The privileges available to each role should be tailorable to each installation.

clinic (see the section on editing lists and entities and the section on system administration below), so that it may always know the default destination of any order: the printer at the front desk, the blood analyzer in the back room, or the clinical laboratory across town.

2.2.1. Order Tests or Procedures

A user with ordering privileges, such as a physician, should be able to order tests and procedures in a variety of ways. The master screen should contain buttons or menu items that take the user directly to the test-ordering module. Other modules, such as the chart module, should also have buttons or menus in appropriate places, to give users quick access to test-ordering functions in situations in which they are most likely to want to order tests.

Within the test-ordering module, users should be able to select individual tests or entire panels of tests from an *ordering menu* of tests and procedures, arranged in a hierarchy of individual tests and named aggregates, or panels. This list of tests, procedures, and panels should be derived from a master list of tests and procedures created for each site. Many physicians, however, may wish to establish their own "short lists" of tests that they order frequently. The workstation should provide editing functions (described below) that allow individual users to compile their own private groupings of tests and procedures to be included in the ordering menu as well.

When a user selects a test or a procedure from the ordering menu, the workstation should present a test-specific form that allows the user to fill out the details of the order. The workstation should fill in as much of the form for the user as possible, to minimize the need to enter repetitive or recurring information. If a clinic always orders a specific prep for a barium enema, for example, the workstation should fill in the default bowel preparation information on the ordering form for a barium enema. Patient-specific annotations may also be appropriate — for example, the physician may wish to indicate special sedation requirements for a bone marrow biopsy or lumbar puncture when the patient is a child. Naturally, the workstation should fill in the name of the ordering physician (the name of the user, or, if the user is not a physician, the name of the physician who has authorized this user to order tests) and the name of the patient. If the user has come to the test-ordering module from the master

screen, the workstation may not know for which patient the tests are being ordered. In this case, the workstation should allow the user to select the patient's name from a menu.[7]

When the user has filled out a test-ordering form, she should be able to close the form and go on to order other tests, prescribe medications, or perform other clinical actions. The workstation should buffer all transactions until the user explicitly signals that they should be carried out. Until that time, the user should be able to review all the actions she has ordered and be able to modify or cancel them. Once the user has instructed the workstation to carry out the transactions, however, they should become part of the medical record and should not be editable. This does not mean that mistakes cannot be corrected; the workstation should provide a cross-out and annotation mechanism that allows users to amend existing records while leaving an audit trail of previously recorded values.

The workstation should make the ordering of tests and procedures as efficient as possible. If a physician is reviewing a patient's white-blood-cell counts in a flowsheet, for example, he should quickly be able to order the tests that will yield new values for the counts. The test-ordering module should be able to note the name and medical record number of the patient, so that this information can be included on the laboratory orders. It should also be able to note that the physician was reviewing the hematology section of the flowsheet, and present him with the list of hematologic tests already selected. This ability to share the context of a user's activities between modules should be a common feature of the workstation's integration.

2.2.2. *Prescribe Medications*

The functions for prescribing medications and therapies should be very similar to those for ordering tests and procedures. The workstation should maintain a database of medications and standard aggregates,

[7] Menus that contain more than a few items can be difficult to use if the only means of navigating among the choices is scrolling, so the workstation should provide alternative ways of navigating menus. Incremental keystroke matching, for example, lets users type the first few characters of a menu choice while the computer automatically scrolls the menu to the first item that matches the characters being typed. If the clinic maintains an active list of one hundred patients, for example, and a physician wants to order a test for Leonard Zagazeta, she might be able to type "zag" when a menu of patients appears to highlight Leonard Zagazeta, and then type a carriage return or click the mouse to select the Zagazeta menu entry.

similar to its database of laboratory tests and procedures, and users should be able to write prescriptions by choosing individual medications or common combinations from menus. The prescriptions that users write should become part of a patient's medical record and should therefore be available to other workstation tools for display in the patient's chart, for printing paper prescriptions, for generating billing information, and so on.

As is discussed in the technical appendices, the workstation's database of medications should contain detailed information about each medication, so that the workstation can offer users varying degrees of support as they write prescriptions. When applicable, for example, the workstation's database should maintain standard dosing information so that the workstation can fill in default values on the prescription forms. As discussed in the section on decision support, below, advanced versions of the workstation should be able to use the knowledge base of medications to provide more powerful assistance, such as warnings of possible drug-drug interactions and the effects of cumulative doses.

2.2.3. Prescribe Chemotherapy or Radiotherapy

The workstation's database of medications should contain information about as many chemotherapeutic agents as possible, and users should be able to prescribe chemotherapies using the same procedure described above for prescribing other medications. However, because chemotherapy is such a central part of oncology practice, the workstation should offer specialized extensions to its prescription functions to manage the ordering of chemotherapies. These extensions should include:

- ***Specialized forms for prescribing chemotherapy combinations***: These are simply aggregates, as described above, tailored to common combinations of drugs (e.g., the combination of procarbazine, vincristine, cyclophosphamide, and CCNU, commonly referred to as POCC). When possible, these specialized forms should contain default information for parameters of drug administration, such as route, pill size, duration, and so on.
- ***Links to textual and graphical data from patient charts***: In the specialized forms, there should be buttons or menus that allow users to review the parts of a patient's chart that pertain especially to treatment. A physician prescribing chemotherapy for a patient should quickly be able to review past administrations, hematology values, psychosocial adjustments to therapy, and past toxic reactions, for example, in either tabular or graphical form.

- ***Links to textual and graphical information:*** As described below in the section on knowledge-access functions, the workstation should provide extensive collections of information pertaining to oncology, with browsing tools that enable users to look up information quickly. These browsers, like the other tools in the workstation, should be invokable from within other components, such as the prescription tools. The workstation should contain mechanisms that enable the browsers to take users directly to the sources of information that are most relevant to the task they were performing when the browser was invoked. For example, a physician using the specialized form for prescribing POCC and wishing to look up information about procarbazine should be able to use a menu or button in the POCC prescription form to invoke the workstation's text browser on any applicable reference work. The prescription tool should be able to provide sufficient information to the text browser that it can immediately display information about procarbazine. These links to browsing tools should give physicians quick access to multiple reference sources containing information about particular agents, combinations, regimens, or protocols. When review of such information is complete, the system should return the physician to the location in the prescribing process from which the inquiry was made.
- ***Tools that offer assistance with knowledge-based tasks, such as calculating doses:*** The workstation should offer a variety of advisory tools, ranging from simple calculators that help users compute dosing percentages, to advanced decision-support programs that can examine a patient's medical record and provide patient-specific advice about dose attenuation, scheduling, and monitoring. (These decision-support programs are discussed later, in the section entitled *receiving decision support*.)

Although the workstation we are describing is targeted primarily at medical oncologists, it should nevertheless offer support in the management of radiotherapy. The workstation should provide forms that allow physicians to record radiotherapy plans and treatments by specifying information such as dosage and port.

2.2.4. Write a Note

In the section on *reviewing a chart*, we described browsers that allow users to review textual portions of a patient's chart. The integrated oncologist's workstation must also provide tools that allow users to generate these texts. Most of the documents in a medical record are

highly structured, so the workstation should provide tools that allow users or system administrators to design templates of the documents they commonly use. These templates can then be used by the workstation's *text editors,* or word processors, to simplify the entry of textual information into the chart. If a physician typically writes his progress notes in SOAP format, for example, he should be able to design a template that reflects the SOAP structure. When he wishes to write a progress note, the workstation can then present him with a template that is already divided into subjective, objective, assessment, and plan sections.

As with other components of the workstation, the text-generating tools should extract from the medical record any information that is always written in a document and include it automatically. For example, if a clinic's progress notes contain a section that includes the protocol the patient is enrolled in, the latest therapy he has received, and his vital signs, the workstation should automatically fill out that section whenever a user creates a new note.

Users should be able to design report templates to include data derived from the medical record, and the workstation's text editors should be able to extract this information from patients' records, just as other workstation components do. For example, if a clinic's progress notes always show the protocol in which the patient is enrolled, his vital signs, and the latest therapy he has received, the workstation should automatically include this information whenever a user creates a new note.

2.2.5. Transcribe a Note

Many physicians will not want to type progress notes or other medical texts themselves; they may prefer to dictate their notes and reports onto an audiotape and hire a transcriptionist to enter the text into the record. An integrated oncologist's workstation should aid in this style of text generation as well. Without additional hardware, the workstation should be able to use its annotation features (described below) to let physicians include markers in a patient's chart that can be referred to in the dictation. For example, a physician reviewing a patient's chart might use a menu or a button to create a new note, but rather than typing text into the empty template, he could annotate it with a numbered dictation marker instead. By speaking this annotation marker at the beginning of his dictation, the physician could indicate where the subsequent dictation should be

transcribed. A transcriptionist listening to the tape could then use the workstation's functions for retrieving annotations to find the marked progress note quickly and transcribe the text into the patient's record.

The design of the workstation should allow for more automated (though more costly) support of dictation as well. Workstation owners should be able to attach specialized cassette recorders or readable-writable optical disc systems to their workstations, and by attaching a microphone to the workstation terminal, they should be able to dictate their reports onto these high-capacity storage devices. The workstation's dictation support functions should be able to annotate these dictations automatically. Later, a transcriptionist could use the workstation's playback functions to retrieve dictations and to jump automatically to the related annotation markers.

It is possible that users will prefer to use other text editors to create progress notes and other reports. The workstation must be able to import text generated by word processors and stored in standard formats and be able to include it in patient records.

2.2.6. Report Complaints, Findings and Diagnoses

Although progress notes present some structure on a printed page or on the screen, they nevertheless consist primarily of text phrases whose interrelatedness can be determined only by a trained human reader. A physician reading a progress note will be able to understand that the patient's complaint of soreness in the mouth, reported in the subjective section of the note, is related (a) to a physician's report of grade 3 oral toxicity, reported in the objective section, (b) to the conclusion that the oral toxicity is due to the high dose of doxorubicin the patient is receiving as part of his treatment, reported in the assessment section, and (c) to the decision to reduce the dose of doxorubicin to alleviate the symptoms, described in the plan section. Computer programs cannot process written natural language sufficiently well to be able to interpret text and to make these sorts of connections from the text alone, yet the information contained in written notes is often of considerable importance to programs that offer decision support. Because the integrated oncologist's workstation will offer decision-support programs that can reason from a patient's data and offer advice on therapy plans and protocol enrollment, it is important that, in addition to offering support for the creation and maintenance of

typed or dictated records, the workstation's patient-record components provide mechanisms that encourage users to record clinical information in a form that makes it accessible to both human and nonhuman readers.

These mechanisms must be designed with care, because busy physicians are unlikely to use a system that makes it difficult or time-consuming to record clinical information in a patient's record. One possible way to provide facilities for semantic annotation is through a mechanism of links. The workstation might maintain lists of common complaints, findings, and diagnoses, and it could make them easily available through buttons or menus. A physician wishing to report a finding, for example, could do so by making just a few selections. At any time, either when reporting the finding or later, the physician could use the problem list tool (described below) to link the finding to other pieces of clinical information. He could associate the finding with a diagnosis, for example, or with other findings that he believes are related to this one, by first selecting each piece of clinical information and then selecting a *link* that describes the association between them. The workstation could maintain a list of common kinds of links, and users could create their own kinds of links as well. Common links might include *causes, is caused by, relieves, is relieved by, co-occurs with, textually explains, ordered to confirm,* and many others.

We wish to emphasize that these facilities for creating semantic links among data elements are not a necessary part of the workstation's design. Past experience with electronic medical record systems has shown that physicians can be strongly resistant to systems that seem to require them to annotate the semantics of the medical data solely for the computer's benefit. At most, facilities such as these should be made available to individual sites on an optional basis.

The functions that permit clinicians to report, and perhaps link, clinical information should be available at many points in the workstation:

- ***In a patient's chart***, so that a physician can record clinical information about that patient
- ***In the test-ordering and prescription components***, so that findings, complaints, and diagnoses can be associated with the clinical actions a physician might take to get more information or to provide treatment

- ***In the note-composition components***, so that the textual renditions of clinical information, which are undecipherable to computer programs, can be linked to the structured renditions of that information

2.2.7. Maintain a Problem List

The functions for reporting and linking clinical information can easily be used to create and maintain problem lists. As described in the section on browsing problem lists, above, the workstation should be able to retrieve information from a patient's record and present it in a format that shows the relationships among the various pieces of information. The problem list format can thus be seen as simply another way of browsing part of the medical record. If, however, an integrated oncology workstation provided facilities for creating semantic annotations, as described above, its problem list could become an interactive component that helps physicians manage the semantic links among pieces of clinical information. When a user asked to see a patient's problem list, the workstation would be able to present a screen that not only let the user browse the information and relationships, but also allowed him to add new relationships and to change or delete old ones. Furthermore, the workstation could display a list of reported pieces of information (e.g., findings) that had not yet been linked to anything, so that users could enhance the cohesion of the medical record at their discretion and leisure.

E.3. Functions for Accessing Knowledge

The integrated oncology workstation should provide facilities that give users access to a variety of information sources. These facilities should range from browsers to knowledge-based expert systems, depending upon what tasks users wish to perform. Many sources of medical information are already available electronically, and some of them, such as PDQ, CANCERLIT, MEDLINE, and digitally encoded textbooks, are quite extensive. Despite the obvious benefit such systems offer to clinicians, however, they are not extensively used. Such lack of use can be attributed to a number of factors: (a) the existence of these systems is still not known to all practitioners; (b) inexperienced computer users often find the systems' interfaces confusing, awkward, or difficult to learn and remember, especially if they consult the programs infrequently; (c) the systems'

use cannot be well integrated into daily clinical routines, either because of problems with physical access (there is no terminal or workstation in the clinic, so the physician must go to the library in a medical center or dial up a service from home), or because of time constraints (most clinicians simply do not have time to go to a computer, turn it on, dial up a remote machine, and perform what may be a frustratingly slow search for information). We believe that an integrated oncology workstation could help decrease the factors that are presently contributing to the less-than-expected use of medical information sources: (1) by providing access to information sources as part of its standard suite of functions, (2) by presenting a uniform and easy-to-use interface to them, and (3) by integrating their use into the routine clinical environment the workstation should create.

3.1. Look Up a Protocol

An oncology workstation should provide access to large databases of clinical protocols, such as that of PDQ. As in PDQ, the integrated oncology workstation should give users a number of ways to search for protocols that interest them:

- ***By identifying the protocol***: Users should be able to specify the name or identifier of a protocol to find a single protocol, or they should be able to specify characteristics common to a collection of protocols: by phase, by drug, by therapy modality, by the organization or provider that offers them, and so on.
- ***By identifying the disease:*** Users should be able to retrieve all protocols that offer treatment for a specified cancer entity.
- ***By identifying the patient:*** Users should be able to identify protocols that pertain to particular ages, sexes, stages of disease, and so on.
- ***By identifying the geographical location of the study:*** Users should be able to constrain the search for clinical trials to those that are offered within a geographical radius.

Users should be able to combine these criteria in a single search if they wish.

The integrated aspect of the oncology workstation should make it particularly easy for users to search protocol databases like PDQs for trials that pertain to a particular patient, as the workstation should be able to extract pertinent information from a patient's record automatically.

3.2. Browse a Protocol Document

The workstation should present the results of its search in a format that allows users to scan a summary of the protocols quickly and to read any section of the protocol in detail. For example, a user might request a search that yields twenty protocols. The user should be able to glance at each protocol's status (active or closed), its sponsoring organization, its research objectives, a synopsis of its eligibility criteria, and other critical information. By selecting an index button or menu, the user should be able to see the full text of all sections of the document describing the protocol. Hypertext links within the protocol documents will further facilitate the browsing of such information.

3.3. Scan for Protocol Eligibility

The workstation should provide functions that automate the process of searching for protocols for which a particular patient may be eligible. This feature could have an extremely beneficial impact on protocol enrollments (the Cancer Information Service of the National Cancer Institute currently handles thousands of telephone calls about this topic every day). Most, if not all, information required to determine eligibility should already be present in a patient's medical record, so a protocol eligibility scanner should be able to take that information and use it to formulate a search automatically. Users should be able to invoke the protocol eligibility scanner any time they wish to scan for therapy options, and they should also be able to specify that the scanner run periodically on its own and report when it finds anything new.

3.4. Perform a Literature Search

The integrated oncology workstation should provide access to many online databases of medical literature through a common interface.[8] Users should be able to perform common searching tasks:

- ***Construct simple or complex queries:*** The workstation should provide tools that help users write syntactically correct search

[8] Note that the common interface the workstation provides may not, in the short term, shield the user from the different searching conventions and terminology used in different databases. In time, the work of the NLM's Unified Medical Language System (UMLS) project, particularly the work on the metathesaurus, may provide a common terminology with which all databases of medical literature may be queried. The first release of the metathesaurus is scheduled for mid-1990, but it will likely be many years before the full impact of UMLS research is reflected in computer interfaces and bibliographic retrieval environments for biomedicine.

formulations. (Some users will still prefer to type their queries free-form, and the workstation should allow them to do so.) Users should be able to specify combinations of topics, authors, journals, and dates in their search formulations by selecting them from menus. The workstation's literature-searching software should maintain dictionaries of search terms (e.g., MeSH terms, names of medical journals, names of major institutions, and others) that users can browse as they formulate their searches.

- ***Refine queries iteratively:*** The workstation should allow users to successively refine a search by editing their formulation and then performing the search again.
- ***Browse search results selectively:*** From the results of a search, users should be able to select groups of citations to examine in greater detail.
- ***Save search results:*** Users should be able to save the results of a search or portions of it to a file, so that they may browse, print, or refine it later.

In addition to these common searching functions, the workstation should be able to link part or all of a search result to a particular patient's record. In this way, workstation users should be able to maintain a record of medical references pertinent to a patient's treatment, or a record of literature searches for billing purposes.

3.5. Browse a Textbook

The integrated oncology workstation should maintain a library of hypertext reference materials online and provide programs that let users browse and search them. Like paper reference works, online references are likely to have different formats, and the workstation's browsing programs should be able to present different formats in a way that makes it easy for readers to use.[9]

If the work is a reference book or another document that has a table of contents and an index, the workstation's browser should allow users both to perform searches on these sections of the document and to scan through them page by page. Selecting an item from the index should use

[9] The format of online reference works may be out of the control of the workstation's designers, and complete browsing functionality may not be fully realizable for all works. The workstation's browsers must be able to handle all major formats, but the technical appendix will discuss the need to establish standards for electronic publishing that support the kind of access we are describing here.

hypertext methods to take the user to the page or pages where the index term is cited; selecting an item from the table of contents should, when the work is divided only into major sections, take the user to the beginning of the section. If the sections of a work are subdivided, the browser should be able to present the table of contents as a flexible outline: users should be able to control the level of detail they see, from a top-level view that shows only the major divisions of the table of contents, through successively detailed views that show the subdivisions of each section. By selecting a heading in the table of contents, a user should be able to instruct the browser to display the subsections under that heading, if any exist, or to display the part of the full document to which the heading refers.

The workstation should provide functions that allow readers of online documents to make their own annotations and hypertext links to other parts of a work, which should be stored by the workstation through the same mechanism that allows users to make links and annotations in a patient chart. These annotations can be in the form of "margin notes" that refer to passages of the text, or "bookmarks" that help a reader return to a passage quickly.

E.4. Functions for Obtaining Decision Support

We have now described how the integrated oncology workstation should maintain data and information about patients in a way that makes them easily accessible to physicians and other health workers. However, as we have described earlier, the patient charts maintained by the workstations will also be used by computer programs that can integrate the information contained in individual patient records with the knowledge stored in specialized knowledge bases to produce a variety of decision support to workstation users. The following sections are organized around some of the kinds of decision support that workstation users might want to receive.

4.1. Obtain Therapy Advice or Critique

In the sections on prescribing medications and chemotherapies, we described a range of tools that helped physicians to prescribe medications and chemotherapeutic drugs. These tools ranged from simple functions that preselected default doses, routes, and frequencies, to complex

decision-support tools that, in keeping with the clinical trial in which the patient was enrolled, suggested the combination of drugs to prescribe and the appropriate tests to order. Advanced advice programs such as these should have the following characteristics:

- ***They should offer patient-specific advice:*** In the section on prescribing chemotherapies, we described programs that provide links between the medical-chart browser and browsers of texts, including the full text of clinical trial protocols. We suggested that these programs should be able to guide the user to the portions of these texts that are most likely to be relevant, given what the user is examining in a patient's chart. Complex clinical advice programs should go further: they should be able to apply knowledge of a specific protocol to a particular patient's clinical data to generate advice that is specific to that patient's condition. The advice should reflect the patient's place in the protocol's time course, his previous response to therapy, the cumulative doses of drugs he has received—in short, the programs should recommend the optimal treatment plan for a particular patient as specified by the protocol. Of course, clinical users should be able to accept, reject, or modify the programs' advice as they see fit.
- ***They should be able to justify their advice:*** Users will often want to know why a program makes a particular suggestion. A physician may want to know, for example, why a therapy advisor does or does not recommend changing the dose of a particular drug. Advice programs might use a number of techniques to explain why they reached the conclusions they did. They could, for example, generate text translations of the reasoning procedures they used to derive their advice,[10] or they could show the user the particular sections of the protocol document in which the rules for determining a particular treatment are found.
- ***They must be unobtrusive:*** Clinical advice programs cannot be designed to engage users in mandatory question–answer dialogs, especially when the users are busy clinicians who have only a few minutes to spend on each case. A clinical advice program must be able to operate while a user performs other tasks, and it should be able to retrieve data already entered into a patient's medical record. When the data it needs are not in the patient's record, an advice program should ask the clinical user for them in

[10] The technical appendices discuss some of the knowledge-representation structures that might be used to encode a knowledge base of protocol information, plus techniques that are commonly used to generate explanations.

an undemanding way. It might suggest by reminders in a dialog box, or by question marks in parts of the patient's chart, that a program the user has invoked needs certain data to complete its task and fulfill the user's request. If the user fails to provide the data it needs, the advice program should be able to reason as much as possible from default information, duly informing the user that it has done so when it offers its advice. When default information is not sufficient, the program should be able to give as much advice as it can and to tell the user what lack of information prevented it from offering more.

Some clinical advice programs should be able to generate treatment plans *de novo;* others should be able to take a physician's treatment plan and *critique* it; that is, they should be able to show, when appropriate, why the physician's treatment plan is a protocol violation or otherwise fails to match the optimal treatment plan. In either case, a user should be able to invoke a recommending or a critiquing program from within other components at any time during the process of recording patient information.

4.2. Receive Other Decision Support

In addition to offering programs that support the creation of therapy plans, the integrated oncology workstation should provide users with systems that support other tasks that require decision-making. In the section discussing modern computing technology, above, we described programs already in use that are able to assist physicians in making diagnoses for a wide array of diseases. In the section describing protocol data-management functions, below, we discuss programs that can scan a database of protocols and suggest which clinical trials for which a particular patient might be eligible. It is not necessary for all these programs to be present on the workstation when it is first introduced; likewise, not every workstation installation site will see a need for all the decision-support tools that will eventually become available. The section on strategic planning discusses these issues of staged and partial introduction of decision-support software in greater detail. Regardless of when or how much decision-support is made available to users, however, all programs should adhere to the same interface guidelines and should integrate with the other programs in the workstation. Whenever possible, decision-support programs should be designed to work fully within the workstation framework; when programs cannot be made to conform to the workstation's style, workstation developers should design front-ends to

these programs that shield workstation users from direct contact with them. In this way workstation users will not have to learn the details of a new program's interface to make use of its services.

E.5. Functions for Protocol Data Management

We have already referred to programs that help determine a patient's eligibility for protocol enrollment, and to programs that help clinicians adhere to the treatment guidelines specified in a protocol. An integrated oncologist's workstation should also offer functions that help physicians, data managers, and other clinical support staff maintain the information required by clinical trials, and it should provide tools that help clinical researchers analyze the recorded information.

5.1. Enroll a Patient

An oncologist's workstation should provide tools that help users enroll patients in clinical trials. It should maintain a database of enrollment-form templates from major groups sponsoring clinical trials and, as with the reports and progress notes described earlier, the enrollment programs should be able to combine these templates with information extracted from a patient's database to produce filled-out enrollment forms automatically. The enrollment programs should be integrated with the protocol advice programs (also described earlier), so that enrolling a patient in a protocol also sets up any information the advice programs might need to begin offering assistance.

5.2. Capture Patient Data

A clinical workstation should provide mechanisms that allow organizations collecting data for clinical trials to capture relevant data directly from a patient's chart. For every protocol it supports, the workstation should maintain a list of the data items being monitored in the clinical trial. It should provide programs that allow data managers to periodically retrieve new data from a patient's chart, as well as updates and revisions to old data. In the absence of a network connection between a clinic and a data collection center, the workstation should provide mechanisms that allow data managers to store captured data on portable storage media, such as floppy disks or portable hard disks. When a data collection center is connected to a clinic via a network, however, the data managers at the

collection center should be able to log on to the clinic's workstation (with appropriate passwords) and to retrieve the data remotely.

5.3. Check Data for Consistency and Completeness

The workstation should include programs that can, using the same knowledge base of protocols the therapy advice program uses, scan the chart of a patient enrolled in a clinical trial and report when data required by the protocol have not been recorded. In addition, these checking programs should be able to report suspicious data values: values that are abnormally high or low, or values that seem to reflect an abnormal fluctuation from earlier values.

5.4. Analyze Patient Data

The workstation should provide suites of tools for analyzing the clinical information contained in its patient records. It should include query tools for extracting data from collections of patient records, and it should provide standard computational tools for performing statistical analyses.

E.6. Functions for Clinic Administration

In the surveys we conducted, physicians reported that administration was one of the fastest growing components of their practices. The integrated oncology workstation must therefore provide tools that assist in this aspect of clinical care.

6.1. Admit a New Patient

The workstation should provide functions that allow administrators, nurses, or receptionists to establish a new patient record in the workstation.

6.2. Export Data to a Billing or Scheduling Program

Many programs already exist to perform billing, scheduling, insurance-reporting, and other administrative tasks, and many clinics have already invested heavily in them. Rather than duplicating these programs, the workstation should, when possible, provide ways of making the data it maintains available to external systems. In time such functions might be incorporated into the workstation itself, but with modern networking

techniques, the large installed base of financial and office management software, and emerging standards for data sharing, we do not believe that a financial/administrative component for the oncology workstation is a high priority feature for early development efforts. As noted in section F, however, most existing programs cannot presently accept data exported from other programs, and the developers of an oncology workstation will have to coordinate their work with the vendors of these systems in order to achieve acceptable degrees of integration.

E.7. Functions for System Administration

The integrated oncology workstation will be a complex system, and it will require ongoing attention and maintenance. Clinics should appoint a person to act as system administrator, and this person should be responsible for maintaining the workstation and for performing workstation-related tasks that fall outside the clinical tasks performed by physicians, nurses, and administrators. This person should not require a technical background or a detailed understanding of the workstation's internal workings; however, some specialized training for the system maintainer will be required.[11]

7.1. Retrieve Remote Data

As much as possible, the workstation should be able to retrieve data from remote information sources, such as laboratory systems, automatically. When this is not feasible, however, the workstation should provide tools that allow system administrators to perform the tasks necessary to retrieve remote data and to integrate them into the workstation's records.

When it is possible for the workstation to retrieve remote data automatically, it should have functions that perform the retrieval and integration without human intervention. However, the workstation should also provide tools that allow system administrators to monitor the exchange of information between the workstation and remote systems so that problems can be spotted more easily. If problems do occur, the workstation should

[11] Vendors of office computing systems have had extensive experience training non-technical staff in clients' offices to be system administrators for networks of computers. Once a system is installed and stable, such people can typically handle routine maintenance tasks in a few minutes per day, calling technicians from the vendor when unusual problems arise.

have tools that help system administrators reconcile them and integrate the transmitted data into workstation's records.

7.2. Custom-Tailor the Workstation

System administrators should be able to custom-tailor the workstation's lists of medications, tests, findings, and relationships to the needs of individual clinics. They must be able to create tables that specify all the tests and procedures available at the site and where each of these procedures is performed, so that the workstation can order them automatically when possible. The workstation's tools should allow system administrators to modify these tables when conditions change (or should propose changes for the administrator's approval based on new information available from the distributed information sources to which it has access). System administrators must also be able to specify the location of the volumes of protocols and electronic textbooks (e.g., whether they are on local optical discs or are accessible through a network connection). As mentioned earlier, the workstation should allow individual sites to design their own printed forms, custom flowsheets, and problem list formats. The workstation must provide tools that help users and system administrators to edit these lists, forms, and tables, and it should provide facilities for storing and maintaining libraries of templates, lists, and tables.

Individual users should be able to modify the lists of medications, tests, and other entities to suit their needs, without disturbing the system versions of these lists. An individual physician, for example, should be able to create her own short list of medications by creating her own private list and filling it with entities drawn from the master list of entities in the system.

E.8. Other Functions

8.1. Read and Write Electronic Mail

The integrated oncology workstation should provide users with access to means of exchanging electronic mail. The workstation should maintain its own electronic mail system, so that all its users can exchange mail with one another. If a workstation is attached to a network that connects it to other workstations, either within the same clinic or remotely, users should be able to exchange mail with users on the other workstations. In

addition, the workstation should support automatic access to other popular sources of electronic mail, such as MCI Mail, CompuServe, Bitnet, and the Internet.

F. Strategies for Implementation

We have described an idealized integrated workstation for clinical oncology, but we fully recognize that many of the individual components will be challenging to develop and to deliver in a cost-effective manner. Such an environment will not be developed overnight, nor will it be possible to introduce the system into clinical settings in a single step. Now that we have described necessary and desirable features of the idealized workstation, we discuss some of the ways to plan for its implementation. It is beyond the scope of this document to provide a detailed design specification or a timeline for development. Instead, we will present some of the strategies that might be used to implement the integrated oncology workstation and some of the issues that must be addressed during its design and development.

Adhere to Software Standards: Designers should adhere carefully to the evolving software standards outlined in Appendix A, so as to maximize portability, flexibility, and independence from specific hardware platforms. Many of the strategies and issues surrounding the creation of the workstation will depend on the choice of hardware, development environments, and other systems-level components to be used in its development. These need not be the components with which the workstation is deployed; because of the falling cost of hardware, the languages, tools, and environments should be selected as much for development efficiency and long-term system maintainability and extendibility as for delivery expedience. If the workstation's developers adhere closely to standards in their choice of basic systems support and in the development of their own modules, the components they create in expensive, feature-laden environments should transfer transparently to less expensive delivery platforms. The most difficult standards choice will be in the area of graphical interface design. Complete standards have yet to emerge for graphical presentations, and it may be some time before satisfactory standards can be achieved. Nevertheless, several good preliminary standards based on X Windows now exist, and implementors should develop the workstation interface within them, while exercising careful software-development practice by creating easily separable modules so that future changes can be accommodated.

Emphasize Integration Early: The need for all the components of the oncologist's workstation to work as an integrated whole should be of primary concern in the early design stages. Specifying the detailed information-sharing needs of each component at an early stage in their design can be a difficult task, but making diversely designed components work well together after the fact is much more difficult still.

Exploit Distributed Computing Technology: In order to achieve the greatest possible modularity, and thereby the largest degree of flexibility in accommodating the diverse setting of oncology practice, the system should be designed using the techniques of modern distributed computing technology. The components of the system, for example, should be implemented within a client/server or an independent cooperating agent model, using carefully designed protocols for communication and control flow.

Incorporate Existing Software: The full range of software required for patient care, record keeping, administration, and protocol research is enormous. In order to minimize the investment in overall software development, the oncologist's workstation should take advantage of existing application software components (e.g., clinic administration, billing, and scheduling programs) as much as possible. Graphical user interfaces and distributed computing can help to weave these components together into an overall system.

Work Within a Network and Communications Infrastructure: The fully functional oncologist's workstation will depend heavily on the availability of information in electronic form and on the ability of resources that allow producers and consumers to communicate electronically. Outside of a few research settings, such resources are almost nonexistent today. The long-term success of the envisioned oncologist's workstation depends on developing this extensive information and communication infrastructure. As discussed in Section III, the creation of such an infrastructure will likely require the cooperative support of federal, state, and local governmental, professional, and commercial organizations.

Identify an Early Target Community: Our survey of oncology practices (Section II.B) indicated that there is a range of practice styles and environmental constraints that preclude developing an initial oncologist's

workstation system that will be usable by everyone. The disseminators of an oncology workstation should therefore not plan to reach all of its broad target community at once. Instead, they should initially identify a smaller sub-community, one characterized by its eagerness to acquire new technology and by the relative ease with which practical and technological barriers to the integration of new technology can be overcome. Our survey suggests that large practices and practices closely tied to hospitals, universities, and medical centers should not be part of the initial target community because of the institutional and logistical difficulties they are likely to face when attempting to acquire new technology. Our survey indicates small-to-medium–sized private oncology clinics constitute the most promising initial market, because they tend to be more autonomous and more anxious to integrate computer technology into their practices. Smaller clinics also tend to have smaller budgets: although it is beyond the scope of this document to prescribe a marketing strategy for an oncology workstation, we suggest that it be priced so that smaller clinics can afford it. The modular design we have described should make it possible for clinics to acquire the full functionality of the workstation incrementally and therefore spread out its cost over time.

F.1. Staged Releases

A plan of incremental investment would also be supported by a strategy of staged workstation releases. A fully functional oncology workstation, such as is outlined in the scenarios and demonstration, entails a tremendous amount of software development. Implementors of an integrated oncology workstation should therefore plan to unfold the workstation's functionality over time, as they achieve major technical milestones that enable them to develop the features in each of the usage categories discussed in section II.E. One approach to identifying staged technical milestones and features is outlined below. We recommend phasing the introduction of the system to provide assistance to the highest priority needs first (flowsheet management, record keeping, incorporation of laboratory and radiology data, assistance with protocol forms filling, and access to literature and protocol databases). This can be followed by the introduction of protocol enrollment and decision-support tools and the full-scale integration of the system with other clinic functions.

1.1. First Release

Key Feature of First Release: the complete support of the electronic medical record, including standard flowsheets, with the capability to handle "results reporting" (automated retrieval of laboratory data from compatible laboratory machines)

Our surveys have shown that the major need for computers in clinical oncology is to support the maintenance of medical records and, in particular, streamlined access to laboratory data. Accordingly, the first release of the oncology workstation should provide an implementation of the key medical record components, including flowsheets. Because automated retrieval of laboratory data is the feature physicians most want to see in the integrated oncology workstation, automated data capture must be present in the first release of the system if the workstation is to be received positively. Providing a standalone version of the workstation first would allow the workstation's vendors to provide a large part of its functionality early: standalone data management, access to information sources via local optical discs, and so on. However, without integrating the workstation with data sources such as laboratory machines, a greater burden of data entry would fall on clerical support staff or even on the physicians themselves. The need to enter data manually, even on an interim basis, could give users the wrong impression of the workstation, and could therefore negatively affect later acceptance of the fully integrated version.

1.1.1. Technical Milestones

a. Development of Database Structures and Functionality

The structure of the patient database and its data dictionary should be designed and implemented as completely as possible by the time of the first release, even though some of their elements will not be used until later releases. Failure to have a completely specified patient database at the outset means future components (such as decision-support systems) that depend on unspecified elements may not be able to function with older patient databases. As discussed in the technical appendix, the patient database should be built within a standard database paradigm.

- *Database Schemas and Data Dictionaries:* The schema for the patient database should be designed to store time-

stamped numerical data and textual data of unspecified length. The data dictionary should contain all terminology used in workstation menus, and it should be designed so that it may be used by components of the workstation to generate menus when needed. The workstation's data dictionary should be able to accept user additions.

- *Storage and Retrieval Functions:* The functions to store and retrieve data in the patient database should be designed with two levels. The lower level functions should implement fast storage and retrieval in the physical database medium, while the higher level functions should be able to create and maintain structures of pointers into the database that allow components to create abstract interpretations, or views, of the underlying data. For example, the low-level functions should be able to store and retrieve time-stamped data, like "200 mg of Procarbazine administered on 1/22/90," while the high-level functions should allow components to identify this administration as part of the patient's third cycle of POCC under a particular protocol.

- *Linking Functions:* As part of the higher level storage and retrieval functionality, the first release should include functions to create and maintain associative links among items in the database. This functionality will be used heavily by the workstation itself in later releases to implement some of the decision-support features, but it will also be required in the first release to allow users to maintain problem lists and to annotate a patient's chart.

b. Development of Mechanisms to Retrieve or Accept Data from Laboratory Machines

- *Network Infrastructure:* The first release of the workstation should contain telecommunications software, built on standard network protocols, that allow the workstation to exchange data with remote machines. Workstation developers should anticipate that most laboratory computers will not be well-equipped to communicate with other computer systems; most of them will have been designed only to print results on a screen or on paper. Although the workstation's telecommunications software should be able to receive data passively, in the early stages of its release it is unlikely that many laboratory machines will have the ability to actively transmit information. The workstation's

telecommunications software must therefore also be designed to poll other machines actively and to retrieve information from them directly. In the first release, these data-exchanging capabilities may consist of little more than the downloading of ASCII text files via terminal-emulation facilities that use dial-out modems on switched telephone lines. Development of these and more sophisticated facilities will require the cooperation of other vendors, with whom the workstation developers will have to work to establish mechanisms for the exchange of data with specific hardware and software systems.

- *Translators:* The integrated oncology workstation should be designed wholly around industry standards for data representation and exchange, and the workstation's developers should monitor the emergence of data-exchange protocols (e.g., HL7) closely. When integrating the workstation with other systems, however, its designers cannot rely on other vendors to adhere to standards. The workstation's data-capturing mechanisms must therefore include translators and generators that can, when necessary, translate all data going in and out of the workstation. At the same time that the workstation's developers are creating mechanisms to contend with nonstandard methods of data exchange, however, they should also seek working relationships with the vendors of other healthcare systems to encourage them to adopt and adhere to industry standards. If the oncology workstation achieves substantial success in the marketplace, the ability to exchange data with it is likely to become an important selling point for other kinds of clinical computing technology.

c. Development of CD-ROM Drivers

The first release of the workstation should include drivers that support access to optical discs in standard formats. Early support of optical discs will not only allow the workstation to provide access to PDQ and to the large quantities of medical literature already published in electronic form, but will also form the basis for mechanisms of software distribution.

d. Development of Interface-Building Tools

Part of the technical effort in the first release should be to acquire tools that streamline the development and maintenance

of the workstation's interface. The importance of choosing an adequate development environment is particularly manifest here, because tools that simplify screen layout and that can integrate display with database retrieval will greatly accelerate the creation of complex user interfaces.

1.1.2. Features of the First Release

Medical Record	- provide standalone medical record support
Knowledge Access	- provide access to existing online knowledge sources and to sources on compatible CD-ROM
Protocol Data Management	- provide libraries of enrollment and reporting form templates from as many clinical trial sponsors as possible, with functions to import data from patient records
Administration Functions	- provide ability to enroll patients into the workstation - provide limited library of administrative form templates
System Level	- provide ability to configure workstation to retrieve data from other compatible machines via dial-out modems - provide operational features that support backup procedures and other reliability measures

1.2. Second Release

Key Features of the Second Release: the introduction of large libraries of medical literature with functions to search, browse, and annotate them; long-distance networks of workstations

1.2.1. Technical Milestones

a. Development of Publication Formats

The formats in which reference materials are currently published electronically may not support some of the workstation's browsing and annotating features. For the second release, the workstation's developers should work with publishers to design or adopt formats that support browsing through document sections and by index and table of contents, full-text searching, graphics, and cross-referencing both within the document and to other materials outside the document (either other reference works or annotations made by readers). When the document formats have been established, the workstation's developers should work with clinical trials research groups to publish as many protocols as possible, planning for a regular schedule of updates and corrections, and with textbook publishers to release important medical reference works in the workstation format. Section III will discuss the issues surrounding the establishment of cooperative relationships with the publishing industry and with sponsors of clinical trials.

b. Development of Text-based Browsers

In step with the design or adoption of publication formats for the medical literature, the workstation's designers should develop the browsing tools that users will employ to read, search, and annotate protocols, text books, hand books, and other reference works.

c. Development of Inter-Workstation Connectivity Functions

The networking capabilities developed for the first release should be enhanced in the second release to allow tighter network communication between workstations. Remote log-in capabilities and remote file-structure access will greatly enhance the ability of workstations to exchange data with each other.

1.2.2. Features of the Second Release

Medical Record

- the ability to create links between medical record, medical literature, and PDQ
- text-based decision support (calculated links between charts and protocol documents)

Knowledge Access	- establishment of distribution procedures for disseminating libraries of protocol documents and other medical literature
	- browsing of protocol documents
Protocol Data Management	- remote access to clinical trials data by authorized data-collection centers
Administration Functions	- the ability to export workstation data to common billing and scheduling programs (This may require coordination with the manufacturers of these programs to get them to accept data from files generated by other programs, rather than solely from their program's interface.)
System Level	- support of long-distance networks of workstations

1.3. Third Release

Key Feature of the Third Release: the introduction of decision-support tools

1.3.1. Technical Milestones

a. Development of Knowledge-Base Formats

The workstation's developers should design knowledge-base structures that capture the knowledge needed by the inference engines, scanners, and classifiers that will provide the workstation's decision-support features. Particular care should be taken to design structures that are easy to create and modify with knowledge-acquisition and editing tools, described below. The design of both knowledge-base structures and knowledge-based programs should be able to support the incremental addition of new knowledge bases and the incremental refinement of existing knowledge bases.

b. Development of Knowledge-Based Programs

The third release should contain a variety of knowledge-based programs whose designs will be based on existing expert-system and classification technology. The technical appendix contains discussions of these subjects.

c. Development of Knowledge-Acquisition and Editing Tools

Crucial to the rapid deployment and controlled maintenance of knowledge-based systems will be the development of programs that aid in the creation and editing of knowledge bases. These systems must be designed for use by non-technical personnel and should be able to generate large portions of knowledge bases without the intervention of technical knowledge engineers. The technical appendix includes further discussion of knowledge-acquisition systems.

1.3.2. Features of the Third Release

Medical Record	- advanced decision support (knowledge-based, rule-based, expert systems providing patient-specific advice)
Knowledge Access	- more libraries of medical literature
Protocol Data Management	- online enrollment in protocols
	- protocol eligibility scanner
	- protocol data checking

F.2. Installation Strategies

The workstation will be installed in clinical settings that will vary considerably in their size, sophistication, and technical resources. The workstation's design should therefore be flexible enough to accommodate site variability without modification to its intrinsic components. As much as possible, the workstation's configuration should be under the control of software "switches" that can modify the system's tables to turn off unneeded functionality (and turn it on again if a site's requirements change) and to establish the appropriate pathways for information access and data exchange. For example, a site that does not participate in any clinical trials and does not have electronic access to laboratory data should be

able to turn off the protocol maintenance functions of the workstation and its data-exchanging capabilities. If the clinic later acquires access to laboratory data by telephone, the workstation's configuration programs should allow a system administrator to establish the protocols for data exchange (from the telephone number of the laboratory's machine to the formats of the data the laboratory machine will provide) and to enable the workstation's data assimilation capabilities. If still later the clinic achieves network access to a new laboratory while retaining telephone connections to the old, the workstation should permit the system administrator to maintain both network configurations.

Because the integrated oncologist's workstation is intended to be used in clinics of different sizes and in oncology research settings, the workstation must be able to accommodate variability in the configuration of users at individual sites. Some sites will have physicians, nurses, and administrators; other sites will have clinical-trial coordinators and data managers. The workstation's configuration software should allow system administrators to define different classes of users, or the *roles* they play. Each role may have a different configuration of privileges and capabilities: for example, the role of clinic administrator may have access to the demographic data stored about a patient but not to sensitive clinical data. A nurse in one clinic may be able to read a patient's chart but not write prescriptions; in another clinic, a nurse may write prescriptions, but only if they are approved by a physician. The workstation should provide some role definitions as defaults, but individual site administrators should be able to modify these to suit the particular needs of their installation.

Section III

Policy and Planning Issues

Contents

A. Impediments to the Vision

Throughout Section II we laid out a vision of what an integrated oncology workstation could provide to the practitioner, but we pointed out in Section II.F that specific strategies are needed for designing a gradual phase-in of functionality. In this section we briefly summarize the ways in which the implementation of the oncology workstation needs to be viewed in the larger context of planning and policy for health information systems in the United States. Computer systems for health care have been evolving for decades, but both system developers and vendors have often been frustrated by the common resistance to computer-based innovations and the difficulty in demonstrating rigorously the value of many of the tools that they have offered. The blame for shortcomings in computer systems has generally been cast at the system developers, often deservedly. However, there is increasing recognition that the utility (and hence the acceptance) of clinical computing systems is closely tied to the incorporation of a critical mass of functionality which simplifies or replaces tasks previously found to be noxious or burdensome. Bringing together this critical mass of functions requires *integration*—a concept we have emphasized in referring to the *integrated* oncology workstation. But, as Section II.E was intended to illustrate, integration implies much more than loading several programs onto the single hard disk of a machine in the oncologist's office. Well-designed interfaces, consistent interactive metaphors, prudent use of distributed data and networking technologies, adoption of standards for linking programs and computers to one another, and incorporation of newer mass storage media such as CD-ROM are all parts of the integration notion that needs to be grasped and adopted. Such features imply not only careful design by system developers but also the creation of standards and an infrastructure which defines the context into which the individual system components must be merged. The development of infrastructure and standards, however, is *not* the task of individual developers; they can only react to what others have put in place. Thus the solutions to many of the impediments we will summarize in this section lie with planners and policy makers rather than with those producing the applications programs themselves.

A.1. Attitudes of Health Workers

Many people have observed that the interest in computing issues among physicians and other health workers has increased dramatically with the introduction of personal computing and, more recently, with the development of well-engineered graphical interfaces. Nonetheless, there are recurring concerns expressed by health workers when they are asked to consider using computers in their daily patient-care work. These attitudes can constitute significant barriers if they are not acknowledged by the developers of the oncologist's workstation and, accordingly, addressed both in the system's design and in the educational materials and training materials that are produced. Typical concerns include the following:

- ***Fear of loss of rapport:*** A recurring concern expressed by physicians is that computer use in their practices will somehow insert the machine in a negative way between the patient and themselves. Such comments reflect the common perception that there is something about computers and their influence that is different from what is observed with the other new technologies (MRI scanners, battery-powered home glucose monitors, and the like) which have been widely embraced by the medical profession. Presumably, this difference has to do in part with the image of the physician personally sitting at the a machine, rather than directing its clinical use. Although there is actually no evidence that the physician's use of a computer will be perceived by patients as anything other than appropriate use of modern methods for information access and data management, it is wise to design systems that encourage but do not require that the computer be used directly by the physician rather than by other workers in the office.
- ***Fear of loss of control:*** Health workers frequently express resistance to the notion that a computer would mandate the management of a particular clinical situation. Many observers believe that this concern is the central barrier limiting the acceptance of computer-based decision-support tools by physicians. It is accordingly clear that advisory systems must emphasize their role as *tools* for knowledge access—performance enhancers that leave the ultimate decisions in the hands of the traditional decision makers.
- ***The problem of inertia:*** For many health workers, the problem is simply that they feel overwhelmed with a hectic lifestyle and with external influences on their practices which they cannot

control. This can make them reluctant to opt for innovations that might also radically change their practice style, especially when the technology being offered is clearly experimental or can be defended only as "holding promise." The importance of immediate payoffs for system users, even with the initial system (see Section II.F) therefore cannot be overemphasized.

- ***Information access versus active decision support:*** Many practitioners will quickly acknowledge their thirst for information access and the potential role of the computer in streamlining that process. They emphasize that the approach must be *easy to learn*, ideally without formal training, and that the tools should be readily accessible at precisely the moment that help is needed. This observation lends further support to the importance of *integration* of information-retrieval resources with other routine data-management functions. However, moving beyond information-retrieval systems (such as CANCERLIT or PDQ) to programs that actively participate in the decision-making process is clearly problematic. As mentioned before, the perception that a computer will tender dogmatic advice leads to reluctance to move beyond simple information-access applications.
- ***Nonacceptance of machine capabilities:*** Another theme that impedes acceptance is a reluctance to acknowledge that computers can reliably assist physicians with complex decision-making tasks in medicine. If a problem is difficult for a physician, then surely the computer cannot be expected to handle it well either! Such attitudes reveal a failure to understand the motivation behind the development of decision-support tools and the realistic limitations that dictate how a computer-based tool should appropriately and effectively be used to help with problem solving. Once again, the challenge for system developers is to address the educational needs that such statements reflect.
- ***Fear of legal liability:*** Another concern , expressed most frequently in the context of decision-support tools, is the fear of legal liability that could accompany their use. Such concerns have clearly constrained not only physician acceptance of decision aids but also the commercial development of such tools. It is difficult for a young company to accept the potential liability for system misuse leading to poor results for specific patients. Without clear legal precedents on the matter, this topic will continue to constrain the development of the health informatics industry.
- ***The challenge of data entry:*** Physicians recognize data entry as being a major barrier to the effective use of computers in clinical

practice. Many practitioners dictate their patient-care notes, often as they move between rooms in a busy practice, and it is accordingly difficult to encourage a model of direct computer interaction to replace the use of the traditional chart. Thus the oncologist's workstation must allow for a range of possible users who will be responsible for patient data-entry. In some practices the oncologists will do this themselves, but in others they will resort to entry by data aids and transcriptionists. Either model must be supported. A well-designed interface, especially one that makes use of modern graphical metaphors for interaction, is more likely to attract the busy physician to hands-on use of the machine.

- ***Reliance on the younger generation***: Among the most common comments from practitioners are those that suggest that current physicians are past the stage of learning how to use gadgetry such as computers and that such technology will be embraced only by members of the younger generation. However, there are remarkably few data to support the notion that current medical-school graduates, once they complete the "indoctrination" of medical training, have attitudes toward computers that are significantly different from those of their mentors. Thus such comments would seem to express a challenge to system developers: "Make it simple and intuitive, like a telephone, and don't expect me to need to know *how* it works in order to *make* it work, and then there is a chance that I will embrace what you have to offer—if it addresses a real need in my practice."

A.2. Costs of Automation

The need for education and careful planning again becomes clear when one considers the role of cost concerns in discouraging physicians from adopting computer-based technologies into their practices. These concerns are often more than monetary. Too many physicians have heard tales of office computer systems that have wreaked havoc with colleagues' practices before the problems were corrected or the systems were finally removed out of frustration. As described in the previous section, threats to provider autonomy may also be perceived as a serious "cost" of adopting computer-based tools into a practice, especially if the machine provides decision-support functions or is perceived as making the clinicians more subject to external monitoring.

The outright costs of equipment and software is of course a consideration, but there is ample evidence that physicians and group practices will

invest heavily in technology when it is perceived as offering them clear advantages in terms of efficiency, quality of care, or the ability to attract patients. The challenge for the oncology workstation, then, is to demonstrate that its costs are justified by the benefits that accrue from its use.

In considering financial costs, the physician is also aware of the start-up expenses associated with converting an existing system. For a large practice, the cost of converting patient records into a new computer-based format could approach the cost of the equipment and software. Carefully designed phase-in strategies are therefore crucial if transition expenses are to be minimized or, at least, spread out over many months.

A.3. Lack of Standards

An impediment that is poorly appreciated by end users but that places a key constraint on system developers is the lack of established standards for data sharing, terminology, and computer-to-computer communication. A variety of groups is addressing the needs for standards (see Appendix A), but none has the weight of the health-care industry solidly behind it. Thus manufacturers and system designers are stymied in their efforts to assure compatibility between new computer-based products and other linked software or hardware which may need to be part of the well-integrated environment. There is a plethora of networking protocols available, for example, any one of which is generally incompatible with the others. Similarly, the data dictionaries used in different medical information systems typically encode local preferences and terms rather than a standardized nosology that is crucial if diverse machines and practices are to share information. If the oncology workstation is to receive lab results from a local chemistry laboratory, for example, it must understand the same terms and reporting units for the tests as does the computer in the laboratory. In the absence of standards, this cannot be guaranteed, even if the two machines can be connected simply via modem and telephone line.

A.4. Absence of Network Planning

We argued earlier that individual system developers cannot be held accountable for the lack of supporting infrastructure, and in no area is this more true than in the field of computer networking. Single users cannot

define the standards, but they can adapt to existing standards that have been defined and broadly accepted. The process of network planning accordingly requires a central coordinating body that understands the diverse needs of the constituents that will need to connect to the network. In a group practice, this means that the total practice needs to define standards with which the oncology workstation will need to comply if it is to gain access to data that are not primarily gathered and maintained in the workstation itself. Similarly, an academic health-science center must coordinate its communications and network plans across the hospital, medical school, and outpatient clinics. On a larger scale, both regional (e.g., state) and national networks require centralized planning and coordination if their full potential is to be realized. In the biomedical community, however, there is a remarkable absence of centralized planning efforts (despite the progress in this area being made in other segments of society). The oncology workstation model we have proposed depends, for its optimal realization, on a coordinated plan for local and nationwide connectivity of computers. Although some databases can be mounted locally on CD-ROM, others will obviously be best accessed via wide-access networks (much as PDQ and CANCERLIT are most commonly accessed today). Coordinated access to all pertinent data will require similar local connections to other computers on which the oncology patient's data are stored. Until the biomedical community recognizes the need for centralized planning in this area and institutes coordinating projects, connectivity between the workstation and other computers is likely to continue to rely on simple modem connections via phone lines.

A.5. Technology Transfer Issues

An obvious question that arises when one considers the oncology workstation model proposed in this document is "Who is going to build and market it?" Despite the large number of companies selling software products to the medical marketplace, most are extremely small and struggling, and the turnover rate in the field is high. None of any size and solvency is primarily involved with marketing to oncologists, and the large computer companies (although several have medical products groups) are reluctant to invest heavily in products where the marketing must largely be done to individual practitioners rather than to large groups. As mentioned earlier, the same legal issues that concern physicians have

constrained the adoption of many computer-related research and development projects by both large and small companies. Furthermore, the medical computing industry has been remarkably slow to adopt the newer concepts of networking and distributed databases which are becoming standards in other parts of the computer world.

The lack of clear-cut mechanisms for technology transfer in a fledgling field has been as serious an impediment to the dissemination of innovation as have been any of the others mentioned in these pages. With risk perceived as high, and many past failures to discourage even the most adventurous, strategies for dissemination that involve shared investment and diluted risk, possibly with government assistance or coordination, are likely to be necessary.

A.6. Trends in Cancer Clinical Trials and Their Coordination

The extension and testing of new basic science findings in the clinic, as well as the identification of optimal ways to use existing therapies, require the meticulous design, performance, and analysis of prospective cancer clinical trials. As our knowledge of the treatment of cancer has increased, cancer clinical trials have become increasingly complex and frequently require the collaborative effort and cooperation of large numbers of geographically dispersed physicians and patients. The complexity of the trials extends to many areas, including the coordination of combined-modality therapies, the enhanced toxicities observed with high-dose-intensity chemotherapy, and the coordination of large, multi-institutional phase III trials.

Paralleling the increased complexity of clinical trials performance is the increased sophistication of the statistical methods used to analyze the clinical trials. This has often increased the amount of data that must be captured on each patient and has consistently required complete and accurate protocol-specific computer databases.

Given the diverse increased demands placed upon those participating in clinical trials, it is not surprising that they have sought methods by which to conduct cancer clinical trials more efficiently. These methods have included the use of microcomputers. Most of the microcomputer systems that have been developed are limited to single-task, single-user

systems that check eligibility, provide tracking of patients on trial, or serve as a database for data analysis. Limited success has been achieved attempting methods of distributed data management. Because large regional and national organizations have been formed to coordinate the implementation of clinical trials, the resulting need for close cooperating and coordination has become both a mandate for enhanced networking with computer-based data capture and, ironically, an impediment to the effective adoption of new technologies. Individual clinics or medical centers cannot make independent decisions regarding the use of electronic databases and computer-based reporting methods; they must look to the central organizations for direction. Thus there is a strong disincentive to automation of clinical trials activities if the regional and national groups do not take a leadership position in this area.

Today the most successful systems are the central general-purpose databases discussed elsewhere in this report, such as PDQ for clinical trials and CANCERLIT or MEDLINE for relevant medical literature. However, none of the computer-based systems has provided comprehensive investigator support, few are closely integrated into the actual clinical setting (for example, as measured by their use in the management of patients not enrolled on clinical trials), and none take full advantage of existing computer technology.

B. Strategies for Addressing the Impediments

As we have just summarized, impediments to adoption of computational methods to assist with the delivery of oncologic care and the execution of cancer clinical trials fall into several categories:[12]

- Attitudes, fears, and entrenched interests
- Legal issues
- Startup, conversion, and operating costs
- Lack of standards and infrastructure, including network planning
- Barriers to commercially viable technology transfer
- Lack of informed leadership

It would be unrealistic to suggest that such problems will yield easily to any single strategy for implementing the vision of the future that we have outlined in this report. If the impediments were largely technical , individual researchers could simply retreat to their laboratories to develop new methods and to implement new prototypes. When the problems are logistical and sociopolitical, however, the solutions will largely await visionary leadership and commitment from those in positions to influence the infrastructure and funding base into which individual systems will be introduced. For the field of cancer therapy, several loci of such leadership will be required:

- Individual clinics and group practices will need to develop a unified commitment to developing common solutions that are well integrated with community resources
- Large healthcare institutions, including the academic health-science centers from which many cancer clinical trials are coordinated, will need to develop a new view of the importance of institutional planning and coordination for both networking and computing services
- Regional health planners, including those who oversee the large regional oncology groups, will need to understand the crucial role they play in defining standards for connectivity and data exchange among constituent practices and hospitals

[12] Some of the material in this section is based on a discussion prepared by Paul Tang and Edward Shortliffe for their work on an Institute of Medicine committee to develop strategies for implementing a forward-looking vision of the computer-based patient record.

- National clinical-research leadership, including the National Cancer Institute and the National Institutes of Health, will need better to understand the role they play in facilitating and providing credibility for efforts to provide coherent standards for communication and data exchange — key issues before the mandatory national infrastructure can be put in place
- A single focus of leadership for planning and coordination of clinical data management is likely to be required. Without a single, credible, and representative organization that can attract cooperation and participation by government agencies, manufacturers, third-party payers, professional societies, and practitioners, it is unlikely that the medical computing field will soon move beyond its current disorganized state (with uncoordinated and incompatible systems the rule rather than the exception).

A committee of the Institute of Medicine, called the Committee to Improve the Patient Record, is currently considering precisely the issues we have outlined here. Their report is due out in early 1991 and it would be premature to predict precisely what their recommendations will be. However, it is clear that whatever else is required, active educational efforts will be mandatory. This report, and the demonstration system we have developed, are examples of the kinds of educational materials that we believe are required if the medical profession, and especially leaders in health planning, are to embrace the vision of coordinated and integrated computer support and are to take steps to ensure that enabling steps are taken so that the concepts we have proposed can be implemented. Only then will the fiscal and logistical costs for individual practitioners drop to a level that will allow the full notion of an integrated oncology workstation to be realized.

Appendix A

Assessment of State-of-the-Art Computer Technology

Contents

A. Introduction

Section II.C presented a high-level overview of the current state of the art in computer technology as it affects the design of an integrated oncology workstation. One of the main purposes of that overview was to make the case, at a functional level, that in order to increase the likelihood that health-care providers will embrace computerized tools such as an oncology workstation, designers must employ state-of-the-art computing technologies (e.g., affordable workstations with high-quality interactive graphical user interfaces, large physical and virtual memory sizes, high processing speed, multiprocessing capability, and distributed network communications facilities). In recommending this type of system design, we are sensitive to the possibility that some readers may believe it to be unrealistic or impractical, but we strongly believe that existing hardware and software technologies are up to the task and that a practical system with capabilities like those sketched in the demonstration prototype (see Appendix C) is technically possible and economically feasible today. Many of the underlying technical details were deferred from the earlier overview discussion. This Appendix attempts to fill in some of those gaps with respect to computer technology.

The presentation is intended primarily for the reader with some background in computing systems. It is beyond the scope of this document, however, to present an extensive tutorial or an encyclopedic account of all the relevant system components. It is important though to understand something about the capabilities of off-the-shelf system components, the range of available technical design options, and what are the major system integration issues facing the implementation of an oncology workstation. The ultimate purpose of this analysis will be to make several points:

- Hardware systems have improved rapidly and dramatically over the past decade in all dimensions of their performance, physical size, reliability, and cost — and this trend is likely to continue at least well into the middle of the 1990s.
- Given the diversity of environments for oncology practices and the diversity of vendors offering computing hardware products, it

is highly unlikely that the oncology community will adopt any standard hardware platform in the foreseeable future.

- Just as in other substantial applications, the major complexity and cost involved in developing, integrating, and supporting an oncologist's workstation will not be in the hardware used but in the *software, information resources,* and *support* required for the various parts of the system.
- Even though early computer workstations have been highly incompatible (e.g., those from IBM, Apple, DEC, SUN, Hewlett Packard, Apollo, AT&T, etc.), software standards are evolving for operating systems, languages, graphical window displays, user interfaces, communications, database systems and query interfaces, etc. that provide significant leverage for developing an oncologist's workstation and facilitate concealing and overcoming intervendor incompatibilities and integrating heterogeneous systems together.
- The most important problem for the design of an oncologist's workstation is to identify a software interface level and distributed architecture on which to base the workstation design — one that will take advantage of the full range of technology needed to facilitate oncologist access but at the same time, take advantage of standards and modularity when implementing integral workstation components so that the system is likely to remain durable, even as technologies like hardware platforms and communications systems advance.

B. Hardware Platforms

In a very simplified view, a typical modern computer workstation includes the components shown in Figure A.1. Without the label "workstation", it would be difficult to distinguish this diagram from those of larger minicomputers and even mainframe computers — and therein lies the power of modern workstations. They are not relegated to be "toy" computers as many of their precursors were, but rather offer a full range of high-performance system services for single or multiple user applications.

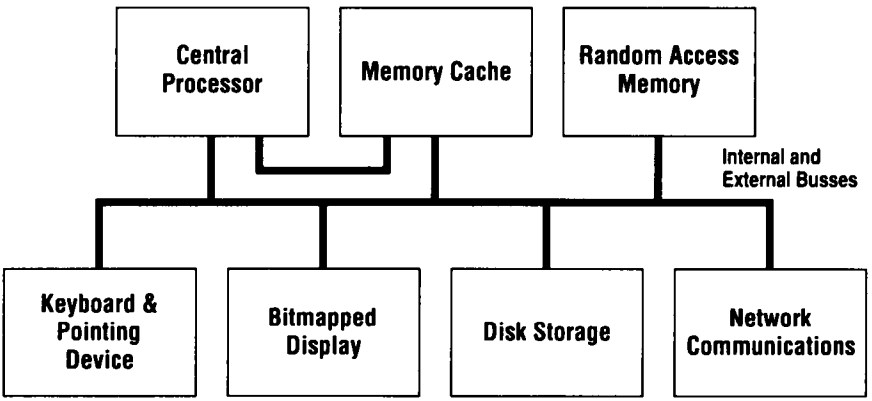

Figure A.1: Diagram of typical workstation elements.

Since the first integrated circuit flip-flops and gates became available commercially in the early 1960s, there have been continuing and dramatic improvements in the density, size, speed, lower power requirements, reliability, cost, complexity, and diversity of microcircuits. Users have become used to seeing a factor of two improvements per year in many of these dimensions. By 1989, microelectronic technology was able to produce chips with millions of interconnected transistors that are able to capture the most complex logic circuits of any sort that humans have been able to design to date — including memories, computer processors, signal processors, communications processors, etc. Industry projections, based on accessible technologies (i.e., without programming in dramatic breakthroughs in areas such as superconducting or optical circuits), anticipate

continued improvements at the same pace until at least the middle of the 1990s when physical limits from the finite speed of light or quantum effects may slow the rate of further progress.

B.1. Central Processors

Modern workstations are all based on fully integrated (single chip) high-performance microprocessors with large (32-bit) data and address path widths, but they vary in the details of their design. The major categories of microprocessors include:[13]

- General-Purpose Processors — These are machines that implement instruction sets comparable to those evolved for mini- and mainframe computers and which are applicable to a wide variety of applications. They are typified by the Motorola 68xxx, INTEL 80x86, SUN SPARC, MIPS R2000/R3000, DEC microVAX, IBM 801, and Hewlett Packard RISC chips that drive workstations marketed by SUN, IBM, Compaq, Apple, Hewlett Packard, DEC, NeXT, AT&T, and many others. These chips are further broken down into:

  - Complex Instruction Set Computers (CISC) — processors implementing a full set of instructions, including complex branching instructions, function/subroutine calling and return with argument passing, memory-to-memory instructions, etc. which require variable amounts of memory to store and numbers of clock cycles to complete. These systems trade simplicity at the software language compiler level for complexity at the hardware level, i.e., a program can accomplish a substantial amount of work by executing a single logical instruction but the underlying hardware has to do a lot of work to fetch the instruction and operands and actually carry out the instruction.

  - Reduced Instruction Set Computers (RISC) — processors implementing a relatively small and simple set of the most frequently executed instructions, such as register to register logic and arithmetic, loads and stores between registers and memory, program branching and flow control, etc. These instructions all require the same amount of memory to store and take a fixed number of clock cycles to complete and have few variables involved so they can be hardwired in the processor for very high speed performance. As a consequence, a program trying to do a more

[13] We consider only single processor workstations in this discussion rather than multicomputers consisting of arrays of microprocessors in various configurations.

complex (but hopefully less frequent) task, such as calling a subroutine, must issue more instructions to do all the steps.

- Special-Purpose Processors — highly tailored processors implementing special functions such as execution of an instruction set idiosyncratic to a particular language (e.g., Lisp Machines); floating point, vector, or other numeric processing; processing of input signals (e.g., speech acquisition); or vector and polygon rendering for 3-dimensional graphics presentations. Such chips are often used as coprocessors with general purpose microprocessors and can speed up the corresponding parts of computation by significant factors.

The jury is still out on which of these approaches will be the most effective or if a hybrid approach will be best. There are high-performance workstations available today using both CISC and RISC technologies (commonly along with numeric coprocessors) that execute at rates of several tens of MIPS (million instructions per second). A key point is that while large processing power machines are cost-effective for use in the oncologist's workstation, it does not make sense to make any part of the system dependent on vendor or system characteristics at the microprocessor level of detail.

B.2. Workstation Memory

Computer programs, the source of the functionality desired for the oncologist's workstation, execute on the central processor out of workstation random access memory (RAM). Memory has been one of the major beneficiaries of microelectronic advances. For a given price, memory chip sizes have doubled each year and continue to improve. Whereas workstations with hundreds of Kilobytes (KB) of memory were commonplace several years ago (e.g., the IBM PC had a 640 KB memory limit and the first Apple Macintosh had 128 KB of memory), workstations with tens of megabytes (MB) of memory are routine today and larger chips are on the way. Memory size is critical for software of the complexity envisioned for the oncologist's workstation simply to accommodate the volume of code involved. Whereas an operating system and early applications ran on small-memory PC's and Macintosh's, more and more sophisticated tools have expanded the needed memory by a factor of 20 or more. At least 2 MB of memory will be required to run the upcoming Release 7 Macintosh operating system and even more memory is needed for running multiple

applications. UNIX-based workstations typically require memory sizes of at least several MB and very large sets of programs can easily swell the requirement to tens of MB.

These needs for larger memories has outstripped the falling prices for memory chips — now on the order of less than $100 per MB for SIMMs (Single In-Line Memory Module) to around $500 per MB for large printed circuit board mounted memories. To avoid excessive memory costs for computers (mainframe computers as well as workstations), a scheme called *virtual memory* (VM) was invented to combine relatively expensive but very high-speed RAM with cheap but very slow secondary disk storage — disk storage is over an order of magnitude cheaper (see below) but is typically one to two orders of magnitude slower. VM works by observing that large computer programs typically reference only a small part of their total address space over time intervals comparable to disk access times for blocks of data. Thus, by breaking high-speed memory into blocks (called *pages*) reflecting the locality of program execution and efficient disk transfers (typically 4 - 32 KB) and with hardware assistance to detect when a program needs to reference pieces of the program not currently in memory, the VM system can allocate physical memory to those parts of the programs that are actually executing (the *working set*) and keep the rest of the program in a *swapping space* on disk. As new memory pages are needed to run a program, memory currently in use is freed up by swapping out pages that have not been active. A well-tuned VM scheme works quite efficiently if the ratio of required virtual memory to physical memory does not exceed about 5 - 10:1.

A related technique, called *caching*, can be used to support the very high-speeds of newer microprocessors. These processor chips are so fast that conventional memory chips cannot keep up with them and so smaller memories of extremely fast chips are introduced to service the most frequent memory requests. For example, if a program is in a loop to accomplish some task, it will repetitively fetch a small group of instructions for every pass through the loop. If these could be kept in very fast memory, the program would run faster. The *memory cache* does this with hardware support to keep data consistent between the cache and main memory when changes (writes) are made.

Long gone are the days when programmer time must be spent worrying about not implementing a desired feature in a program because it will take

too much memory space or be spent managing complex program overlays to get around memory limitations. Modern systems provide a hierarchy of economical memory resources that can accommodate programs requiring hundreds of MB of memory. This technology must be exploited in the development of an oncologist's workstation.

B.3. Mass Disk Storage

An oncologist's workstation is data-intensive and requires access to randomly accessible mass storage for general file space, patient record database elements, knowledge base elements, or virtual memory swapping. Rotating disk technologies of various sorts have historically met these needs and continue to do so.[14] While cost and performance measures for disk technologies have not improved quite as rapidly as microelectronic technologies (disks are basically mechanical devices), impressive gains have been made nonetheless. The primary alternatives for disk support include:

- Removable Medium Magnetic Disks — the principal options are 3.5 inch or 5.25 inch floppy disk drives. These disks store 1-2 MB of data and cost about $100 per MB.
- Fixed Medium Magnetic Disks — magnetic disk technology is well established and devices are available for workstations with storage capacities ranging from about 40 MB up to 1 gigabyte (GB), all with a form-factor that allows installation inside of desktop workstation configurations. These disks have seek access times of 10 - 20 msec and costs ranging from $5 - 10 per MB.
- Removable Medium Optical Discs — the technology for optical disc storage has been available for a number of years but standards have been slow to get established. Motivated largely by the commercial success of compact discs for audio recordings, digital optical devices are now more widely used. Optical disc systems are available that are *read-only* (CD-ROMs), *write-once-read-many* (WORMs), and fully read-write (like magnetic disks). The most common optical cartridge size has been 5.25 inches, with storage capacities of 300 - 400 MB per side. Recently, 3.5 inch cartridges have come into use with capacities of about 130 MB. Because of the more complex recording and playback mecha-

[14] High-volume magnetic tape technologies are also available but because of the serial nature of tape systems, they are not considered as appropriate for random access applications. They do have a key role for file system backup and file archiving.

nisms, optical discs are typically about 5 times slower in seek access than magnetic disks. Cartridge changers (or jukeboxes) are available in various sizes to provide a total on-line optical storage capacity of tens to hundreds of GB. A typical cartridge swap time for jukebox mechanisms is about 10 sec. Large-volume Read/ Write optical storage systems cost on the order of $1 - 2 per MB.

C. Human-Computer Interfaces

A crucial part of any workstation that is to be used by physicians and other health care professionals is the human-computer interface. The invention of interactive computing in the early 1960s, based initially on the "typewriter" interface metaphor, brought an immediate need to develop better ways to present information to users and for users to communicate with programs. As analog and digital hardware technologies have evolved, various devices taking advantage of visual, audio, and tactile sensory modalities have been tried. While far from being a solved problem, the most important advance in workstation interface capabilities has been the development in the early 1980s of the integrated Macintosh-style interface by Apple Computer, based on the pioneering research work on graphical user interfaces (GUI) at Xerox Palo Alto Research Center during the 1970s. If copying is the sincerest form of flattery, then perhaps the best indication of the power and impact of this approach is the dramatic rush by all parts of the computing industry to provide and use Mac-like interfaces on many of their workstations.

The Macintosh interface is an integration of conceptually simple hardware technologies with a standardized set of software tools, graphical *windows*, and interface *look and feel* guidelines that provide a much more intuitive way for users to visualize and control what is going on in the computer. The core hardware consists of the familiar keyboard, a bitmapped television screen (a rectangular raster of dots, each of which can be turned on or off in the simplest kind of system or assigned a shade of gray or a color in more complex systems), and a mouse pointing device (which causes a visual mark on the screen to move in response to movement of the mouse device on a desktop and has one or more buttons that can be pressed to send simple signals to the computer). Other kinds of pointing devices sometimes used include light pens, track balls, or touch-sensitive screens. This hand-eye interface is augmented for some applications by computer-generated sound output. The enabling hardware developments come from those we have already discussed — a) memories large enough to store and manipulate the information needed to paint large screens and b) microcomputers powerful enough to process all of

the interactive events of keyboard strokes, mouse movements, and display screen changes seemingly in real time as application programs execute. It is now commonplace for workstations to have screens that can display more than a million distinguishable picture elements without any sense of screen flicker. The controlled shading of each of these picture elements (by the computer) produces an image that can represent arbitrary text, drawings, or pictures which in turn can be interpreted and manipulated by the user.

This much of the interface technology, that is necessary and sufficient for the initial oncologist's workstation, exists now and is readily and economically available on standard, commercial workstations. Future systems may include less bulky liquid crystal displays (LCDs) with raster sizes comparable to today's television-type CRTs — some systems with LCD displays are already available, e.g., the Apple portable Macintosh and the GRID computer. These could be used to develop "clipboard" workstations that a physician can more easily carry around and use in a clinical or hospital setting and that, with a stylus (light pen-like) device, could allow handwritten information entry and system control (several companies, e.g., Apple, IBM, and Go, are already reported to be working on such stylus computer systems). Similarly, future systems may feature discrete or continuous speech recognition interfaces — several discrete speech systems are already on the market for specialized applications such as preparing diagnostic radiology reports and a number of more sophisticated continuous speech systems are available for research work.

D. Communications Systems and Networks

Just as the telephone, FAX, and US mail systems facilitate conventional interactions between people, computer-based communications systems facilitate interactions between computer users and access to remote computing and information resources, e.g., a physician connecting to review recent laboratory results or wishing to search the National Cancer Institute PDQ database for therapy protocol advice relevant to a particular patient. With the development of interactive computing in the early 1960s, communications between terminals and computing resources became a high-priority target for innovation. Since modern workstations used in a distributed computing environment may simultaneously be in contact with each other, file servers and databases, printers, other computer resources, instruments, etc., the need for flexible and high-performance intercomputer communications is re-emphasized. Most intercomputer communications depend on various methods for the serial (bit by bit) transmission of information (as opposed to very high bandwidth parallel interconnects used internally in computers, e.g., for backplane busses). These serial communication technologies are highly developed now with many commercial offerings available for equipment and software:

- Hardwire point-to-point connections — these links provide physically fixed connections between terminals, workstations, and other resources (computers, printers, instruments, or network interconnection devices — see below). They are highly variable in length (relatively short copper wires to long dedicated telephone circuits) and are best used for permanent installations. Depending on the kind of hardwire involved (twisted pair, baseband or wideband coaxial cable, fiber optic cable, etc.), low to very high speeds are possible.
- Telephone modem connections — these links take advantage of the switching facilities of the voice telephone system to allow a workstation to connect to any other device with a corresponding modem interface as needed. As in the case of FAX connections, modem links are initiated by dialing a destination telephone number and negotiating a method for encoding and transmitting data. Subsequently, data is sent back and forth until the session is completed and then the connection is terminated. This method

offers high flexibility, relatively low cost for only intermittently needed connections, and moderate speeds (9600 bit/second modems are commonly in use today on the switched telephone network).

- Local Area Networks (LANs) — an easy way to think of local area networks is that they are similar to telephone party lines. Several devices are connected to a common transmission line and they coordinate among themselves as to which will transmit at a given moment. Data is sent in the form of packets of information, each containing appropriate source and destination addresses, control and error detection codes, and the data itself. Protocols and rules of etiquette are established so that communication is efficient, any device can select any other device to communicate with, collisions between two devices attempting to transmit at the same time can be resolved, and one device does not dominate the resource if others need to communicate. Of course, computer LANs entail the same sorts of privacy and security issues that party lines do in that any of the connected devices can listen to the conversation between any other devices. Numerous ways of organizing LANs are in use, including Ethernet (branched bus systems), token rings, stars, etc. Bridges and gateways provide ways to interconnect LANs with each other and with other communication systems. Individual LAN segments are typically less than 1 km in length and, depending on whether twisted pair, coaxial cable, or fiber optic cable is used as the transmission medium, can transmit at speeds from several Mbits/sec to 100 Mbits/sec.

- Wide Area Networks (WANs) — over very long distances, the coordinating protocols used for LANs ("party line" networks) become inefficient and cease to work. Like LANs, WANs use transmission lines shared among multiple users to transmit packets of information but, instead of sending each packet directly from source host to destination host, WANs use dedicated network computer nodes and store-and-forward techniques to move packets progressively between the source and destination hosts. LANs are connected to larger WANs through gateway machines that handle the buffering of packets into the larger internet (network of networks). The lines interconnecting WAN nodes can use most any available medium and transmission mechanism — many protocol systems have been worked out.

  Perhaps the best known WAN is the National Research and Education Network operated by the National Science Foundation.

The current NSFnet includes a backbone of T1 lines (1.5 Mbits/sec) spanning the US and connecting to Regional networks that provide access for many academic campuses and research laboratories. The NSFnet backbone will be upgraded to use T3 lines (45 Mbits/sec) at some unspecified future time.

While many public and private protocols have been developed for intermachine communications, the most widely implemented and used one is TCP/IP (Transmission Control Protocol/Internet Protocol) and a rich suite of higher level service protocols that have been defined on top of it (e.g., TELNET for terminal-to-host communications, FTP for file transfers, SMTP for electronic mail communication, RPC for remote procedure calls, etc.). In the long term, these will be replaced by a set of internationally agreed upon protocols defined by the International Standards Organization (ISO) but these definitions are currently incomplete and few implementations exist.

E. Systems Software

Systems software layers provide the interface with computer hardware (processor, memory, and peripheral devices) and turn these resources into useful services for supporting user information processing activities. These services connect keystrokes on the keyboard with running programs, support file naming and access, provide network services like EMail and file transfer, implement language processors for creating new programs, support graphics interactions, etc. In a simplified view shown in Figure A.2, there are three layers: the operating system, utilities and services, and application programs.

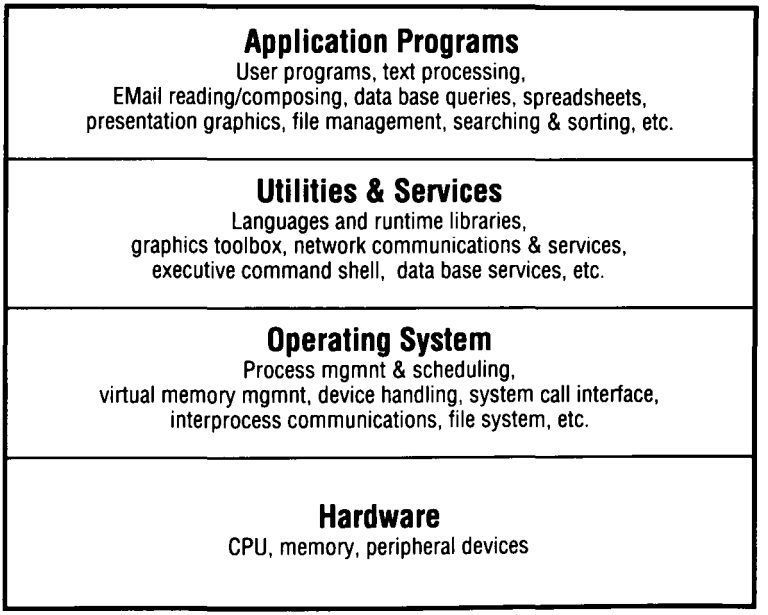

Figure A.2: Diagram of systems software layers.

In the following sections, we summarize some of the diverse services and tools provided by systems software.

E.1. Operating Systems

Every computer or workstation has an operating system (OS) which begins to execute as soon as the machine is turned on or rebooted. Some of the more common workstation operating systems include UNIX which runs on a wide range of computers (sometimes with slightly different names), MS/DOS and OS/2 which runs on IBM PCs, VMS which runs on DEC workstations, and Macintosh OS which runs on Apple workstations. At the lowest level, the OS handles the various pieces of hardware that comprise the workstation — microprocessor, memory, keyboard, mouse, display, disks, printer, communications links, etc. Each different device has a particular way of communicating data to and from the central processing unit (CPU), has to be handled in an appropriate sequence of steps, and has certain error conditions that may arise. The OS device drivers, one for each device, take care of these tasks.

At a higher level, the OS "repackages" the low-level device capabilities into services useful in processing information. For example, a magnetic disk is basically a spinning platter with a read/write head that can index radially in discrete steps to transfer bands or tracks of raw binary data between the CPU and the disk surface. The OS file service turns this raw data storage and retrieval capability into the ability to define named user directories and files of various sorts that have appropriate access control and security and can be manipulated by logically meaningful commands — open, close, read, write, append, copy, etc. The file system also provides an essential set of tools for the system itself in storing and making accessible the many programs that the system uses to provide still other utilities and services.

Another essential service provided by operating systems is CPU and memory management. A raw CPU basically fetches and executes instructions sequentially from memory until an instruction commands a branch to begin fetching and executing instructions from some other memory location or an outside device event (interrupt) occurs that transfers control — the machine has no way of knowing that it is executing a text processing program or a spreadsheet program or the user command executive program. How does the system make sure the proper programs are loaded in the proper part of memory at the proper time and that the system is consistently set up and controlled to safely carry out arbitrary

user commands? The basic system service that makes this possible is the *process*, a logical unit of executable program which is allocated memory and processing resources to carry out its task. For example, when a user asks that a file be copied from one place to another, the system looks for the file-copy program, figures out how much memory it needs to run, allocates that amount of memory from the pool of unused memory, loads the file-copy program into it, runs the program until completed, and then returns the memory to the pool. The combination of program, allocated memory, and any other resources necessary to run the program (such as stack space, register set, etc.) is called a process. The operating system's ability to create, run, control, and delete arbitrary processes flexibly provides a powerful means of servicing user information processing requests. Systems that can handle one process at a time are called uniprocessing systems (e.g., MS/DOS and Macintosh OS), and those that can handle more than one at a time are called multiprocessing systems (e.g., UNIX, VMS, and OS/2).

Multiprocessing systems take advantage of the fact that in interactive computing, a computer is often fast enough that it can finish one task before the human user can request the next one. For example, while entering text in a text editor on a 10 million instruction per second (MIPS) workstation, the time between typing successive characters (perhaps a fifth of a second) is enough time to execute two million instructions doing some other task! These resources can be used to do work for the same user (fetching a remote file or sending a previously composed EMail message) or to do work for some other user. In systems like the oncologist's workstation, where a number of things may be going on simultaneously, e.g., entering information or progress notes, fetching new information from the patient record, assessing the appropriateness of a therapy plan, etc., a multiprocessing OS is essential.

Other kinds of services provided by modern operating systems include virtual memory handling, process scheduling based on assigned priorities or the interactive nature of the task each process is doing, mechanisms for processes to communicate with each other, and a system call interface so that processes can themselves ask the OS to perform services for them (e.g., file access, creating still other processes, input/output, interprocess communication, etc.).

E.2. Systems Utilities and Tools

A rich collection of utilities and tools have been built on top of the above OS facilities that allow the user to control what the systems does for them and to create still more new programs to extend the capabilities of the system. The range of services available is much too extensive to enumerate here — large manual sets are devoted to this task — but we will give a sampling of some of the more important examples. The basic interface between users and an interactive computing system is defined by the *Executive* or *Shell* program. When one establishes a new connection to a system and is greeted by a "Please log in" message, the system has noticed activity on a previously unused connection port and has created a process in response that runs an executive program. On a stand-alone workstation, when rebooted, the system automatically starts up an executive program and waits for a user to interact. The executive program may be command-based (as is the CShell on UNIX) or direct-manipulation graphics-based (as is the MultiFinder on the Macintosh OS), but each provides a wide range of basic services for customizing the user environment, navigating through the file system and manipulating files (locally or on another machine), executing and controlling application programs, establishing network connections and manipulating them, obtaining system information, doing input/output operations such as displaying or printing a document, etc.

On multiprocessing systems, many background services are typically in operation to keep the user apprised of events of interest (e.g., the arrival of new mail) or to carry out tasks (such as printing a file to a busy printer) while the user moves on to do other things. Similar daemons facilitate outside network connections for terminals, file access, EMail transmission, network routing information updates, remote interprocess communication, etc. Resource directory databases are maintained so that the user can access services by name rather than having to know where they are physically on the network.

When the user runs an application program (e.g., text editor, spreadsheet, or EMail reader), that application in turn depends on a large set of program development tools and shared system runtime libraries and services that make its function possible. The most obvious of these are the programming languages used to implement applications. Some of the more common ones today include C, FORTRAN, PASCAL, and CommonLisp with object-

oriented extensions to them (e.g., C++, Objective C, and CLOS). Each of these languages has associated with it an interpreter and/or a compiler, debugging tools (often including data structure inspectors, execution control like single-stepping, and incremental program change tools), and extensive function libraries for operations like calculating statistical or mathematical functions, string manipulations, controlling network connections, handling files, database query access (e.g., functions for creating and processing SQL queries to database servers), input/output operations, and managing graphics displays, etc. Because of the increasing importance of graphical user interfaces and the difficulty of implementing them using classical textual programming techniques, new direct manipulation graphical tools are being developed to assist in this process. Some current examples include NextStep, Action!, HyperCard, and ProGraph.

E.3. Applications

The applications layer comprises the programs that users use to carry out their work. A wide variety of tools are familiar to assist with text editing and document preparation, drawing and presentation graphics, EMail reading and composition, business processing and spreadsheet systems, statistical analyses, database querying, and many other functions. The oncologist's workstation software will run as a set of applications in the context of other existing tools that will complement its patient care and administrative services. In some system environments, notably UNIX, applications have been developed by many different people with few uniform conventions for names or for the style of user interface commands. These systems, while very powerful, tend to be very difficult for the average user to master because there is so much to remember and little pattern to guide the process. Other systems, beginning with the Apple Macintosh, have defined a standard set of *look and feel* conventions that are applied across all applications. These standards make the interfaces to diverse applications similar and relatively easy for the user to figure out, even when the applications deal with very different tasks. There is a standard way to select and open files, a standard way to perform editing functions like cutting and pasting, a standard way to save files and print output, etc. A successful implementation of the oncologist's workstation software must take advantage of conventions such as these for all of the diverse operations involved to facilitate their integration and ease of use.

F. Standards

Historically, the high cost and centralized nature of computer equipment forced a strong commitment to one vendor or another and compatibility issues revolved around whether new releases of the system software or upgrades to the hardware from that vendor would be upward compatible. It was not expected that significant applications written for, say, an IBM mainframe would run on a DEC VAX without substantial modification, even when the program was written in a standard language like FORTRAN or COBAL.

With the rapid evolution of inexpensive microprocessors, workstation computers, and distributed computing, however, the cost of hardware has plummeted and the primary cost is now in the software developed for user applications. This, in turn, has led to the need for broader standards to ensure software portability, interoperability, and information communications among diverse systems in order to preserve the growing investment in software. Organizations including user groups, professional societies, industry consortia, and government agencies are playing important roles in evolving standards at all system levels. Since most computer and workstation vendors view themselves primarily as being in the hardware business (i.e., software is an important but secondary business activity) and differences in their respective hardware designs are what differentiate one vendor from another, it is not likely that any standard hardware design will emerge for computer workstation components or systems. The primary standardization at the hardware level is in terms of the backplane busses used to interconnect processors, memory, and high speed devices and the interfaces for peripheral equipment. We do not expect compatible hardware details for workstation designs across vendors any time in the near future.

Thus, in order to preserve the usability of complex and expensive-to-develop software as new hardware emerges, we can try to take advantage of the layering of systems software discussed earlier to create a compatible software interface across diverse hardware. This is an old idea, of course, in that computer languages such as FORTRAN were standardized quickly, after they were invented and it was recognized how powerful they were for developing applications software. Thus, if a program is written in

FORTRAN and there are FORTRAN compilers for two different machines, one would like to expect that a program written on one could be moved to the other, recompiled, and run without problem. Unfortunately, language standardization is only part of the picture. Differing file name conventions between systems, graphical interface systems, data word sizes, operating system interfaces, communications services, etc. confound the transfer. Thus, software standardization at all levels of the system is needed.

Fortunately, these standards are evolving slowly, under pressure from diverse groups, including some computer vendor companies. Some of the principal groups involved include the federal government, the Institute of Electrical and Electronic Engineers (IEEE), the American National Standards Institute (ANSI), the International Standards Organization (ISO), the Open Software Foundation (OSF), and UNIX International (UI). We cannot detail the status of all of the standards involved — indeed, the negotiation of many of them is fraught with deep technical and non-technical difficulties and delays — but we list the principal ones relevant to the oncologist's workstation implementation:

- Operating System: UNIX representing a merger of University of California 4.3 UNIX (developed with support from the Defense Advanced Research Projects Agency — DARPA), American Telephone and Telegraph System V UNIX, and Carnegie Mellon University Mach/UNIX. All other frequently used operating systems (e.g., DEC VMS, Macintosh OS, and IBM OS/2) are highly vendor-specific.
- Network Communications: the DARPA-developed Transmission Control Protocol/Internet Protocol (TCP/IP) and its subsidiary information exchange protocols (e.g., terminal connections — TELNET, file transfer protocol — FTP, simple mail transfer protocol — SMTP, etc.)
- Languages: C and CommonLisp, including their object-oriented programming extensions, C++ and CLOS. FORTRAN is used widely for numeric and statistical computations and COBAL is used for many business applications.
- Windows and Graphics: the X-Window display client/server system with toolbox library implementations accessible from the various programming languages, developed under the MIT project Athena in collaboration with Digital Equipment Corporation. The actual appearance style of the displayed objects (windows,

menus, scroll bars, etc.) is referred to as the look-and-feel and standards for these are still evolving. At present two main UNIX-related candidates exist — OSF Motif and UI OpenLook.

- Database Queries: Relational database systems with Standard Query Language (SQL) interfaces.

G. Integrated System Architectures

The range of workstation hardware, software, and communication systems described above allow for many diverse architectures in allocating computing resources and services between systems for various interactive computing applications. The oldest configuration is a terminal connected to a computer, often, in the past, over long distances because of the scarcity and high cost of adequate computing devices. Even the modern, self-contained, desktop workstation is built based on this configuration but with a very short communication link between the terminal and the microprocessor. Such workstations are often used themselves as clients to remote computing services or to communicate with other workstations. This client/server model provides a highly modular and flexible way of organizing computational tasks which can be distributed among diverse systems in an almost endless variety of ways. A user can be writing a scientific paper using a graphics-oriented word processor on their local workstation while connecting to the National Library of Medicine (or a local CD-ROM server) to do a MEDLINE bibliographic search, copying and reformatting a result into the bibliography of the paper, and sending the draft off to a shared network laser printer server to get a hardcopy without any concern for the details of how the communication and processing is done. These powerful tools form the basis for our conceptual design of the oncologist's workstation, providing the necessary user-friendly access to the many information and decision-making resources needed for effective oncology care.

Appendix B
Architectures for Implementation

Contents

In support of Section II.F, this Appendix summarizes the elements of the oncologist's workstation and discusses some of the technical and strategic issues underlying its implementation and deployment . This analysis is based on the technological background discussed in Appendix A.

A. Elements of the Oncologist's Workstation and Current Practice

The oncologist's workstation described in the main body of this report and illustrated in the SuperCard-based video tape demonstration (see Appendix C) must be integrated deeply into the clinical environment for which it is intended as well as into that of the oncology protocol study groups. The workstation includes support for clinical patient care, clinic administration, and study group activities:

- Patient care
  - Patient data collection and medical record keeping
  - Decision making and therapy planning and execution
  - Progress notes, communication, and correspondence
  - Information resource access
- Administration and business
  - Patient and clinic scheduling
  - Pharmacy and clinic inventory control
  - Insurance, billing, and financial management
  - Business planning
- Protocol study support
  - Maintaining protocol/regimen knowledge base
  - Physician protocol management support
  - Data transfer and assembly
  - Statistical analysis and interpretation

To date, based on our survey of oncology out-patient clinics reported in Section II.B, the predominant use of computers is in clinic administration and business management, in the recording and reporting of data in some clinical laboratories, and in the statistical analysis of clinical trial data by study teams. There is almost no use of computer tools in patient care, only very spotty use of on-line information resources such as MEDLINE and PDQ, and almost uniformly manual filling out and handling of protocol eligibility and data reporting forms.

From a technological point of view, this situation presents both significant opportunities and daunting challenges. On the one hand, there is lots of room for computerized patient record and protocol management tools to increase efficiency and save time for physicians and other health care providers, thereby increasing the likelihood of protocol use. On the other hand, each of the clinics we investigated was almost entirely different from the others with highly idiosyncratic methods and procedures for managing patient care, patient records, the acquisition of outside data from laboratories, administrative functions, participating in and reporting protocol data, and allocating responsibilities between physicians and other clinic personnel.

We have emphasized that the ideal way to introduce the use of computerized tools into oncology care is to do so by emulating the ongoing office practice as much as possible so as to ease the transition problems of training and procedure disruption while installing the system. However, the diversity of modern oncology practice represents an incredibly heterogeneous environment for the design and integration of such a system. Because of the lack of standards or common procedures in current practices, each system would have to be uniquely designed and configured, raising the attendant implementation and support costs to unreasonable levels. Thus, the challenge in formulating an implementation strategy for the oncologist's workstation is to design a system with as much flexibility and adaptability as can be achieved cost effectively and to phase its introduction first into clinical practices that are motivated and able to adapt their procedures over time to the new tools.

B. Alternative Architectures

The computing technology described in some detail in Appendix A offers significant advantage in dealing with the current diversity of oncology practices, including the frequently existing previous investments in administrative computing systems or services and the cumbersome interfaces with external laboratory and record systems:

- The falling costs of high-performance hardware systems means that the overall system can include multiple machines, each potentially running software tuned to a different phase of clinic function. This flexibility means also that parts of the system can be upgraded to increase capacity or improve cost effectiveness without affecting other parts of the system.
- Distributed computing technology, using the client-server model of system organization in conjunction with multitasking operating systems, intercomputer communications, distributed file and relational database systems, etc. provides the flexibility to link heterogeneous hardware and software systems together effectively.
- Graphical User Interfaces have been proven to provide a more intuitive, convenient, and effective way for diverse system users to communicate with and control workstation functions.
- The availability of cost effective large address spaces and virtual memory systems means that programmers can optimize workstation software for user convenience and support as well as using design techniques that incorporate a maximum flexibility. The languages, tools, and environments should be selected as much for development efficiency and long-term system maintainability and extendibility as for delivery expedience.
- The emergence of system software and communication standards means that the oncologist's workstation can be implemented so as to preserve the investment in software development, even as hardware technology continues to change rapidly. The most difficult standards choice will be the graphical interface design, as complete standards have yet to emerge for graphical presentations and may take a long time to be satisfactorily standardized. Nevertheless, good X-Windows-based choices exist for these parts of the system and they can be carefully separated into identifiable modules so that future changes are easier.

These technological considerations allow for a range of alternative models for designing and constructing the oncologist's workstation as shown in Figures B.1 through B.2. Figure B.1 illustrates a "classical" design in which there is a central server in the clinic which provides all computing, file storage, record keeping support, printing, communications, and other services to users who have access to the system through terminals. By using X-Terminals for users, the environment can take advantage of powerful graphical user interfaces and share a higher performance (and hence costlier) server machine. X-Terminal hardware consists of a bitmapped display screen, a keyboard and mouse, a local processor with a relatively small amount of memory, and a network connection. This is coupled with software that allows it to act as an X-Window display server. An X-Terminal does not have enough resources to perform other computing tasks as is commonly done on a personal computer workstation. Thus, an X-Terminal (as opposed to older character-oriented terminals) is an inexpensive way to support a graphical user interface for (X-Windows-based) applications running on a remote machine.

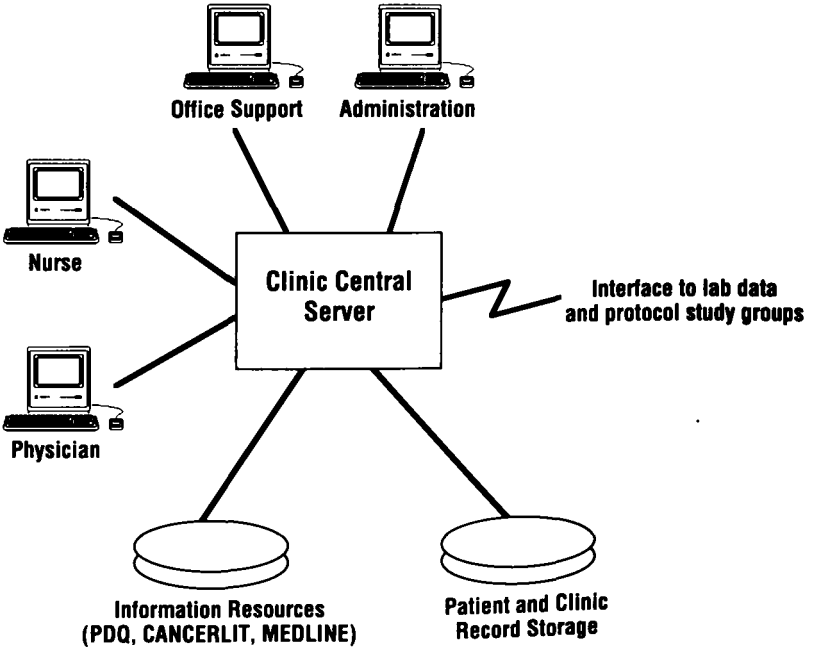

Figure B.1: Centralized Server Model

The central model requires that all software used by clinic personnel run on the same computer system. If parts of the system that require substantial computing resources are infrequently used by clinic staff, this model gives a way of sharing the cost of these services over the entire staff since each person uses a common central server.

Figure B.2 illustrates a distributed computing model, using local servers for file storage, shared computing, and remote communications such as to laboratory systems or protocol study groups. In this model, each user has a workstation on their desk which provides a graphical user interface and independent computing resources for frequently used services that are not subject to economies of scale. This model does not require that all of the workstations be identical and so a physician may use a different type of system than the clinic administrator. This allows a tailoring of the computing environment to support diverse pieces of software that might be available and also to increase the capacity of the system by adding workstations as needed.

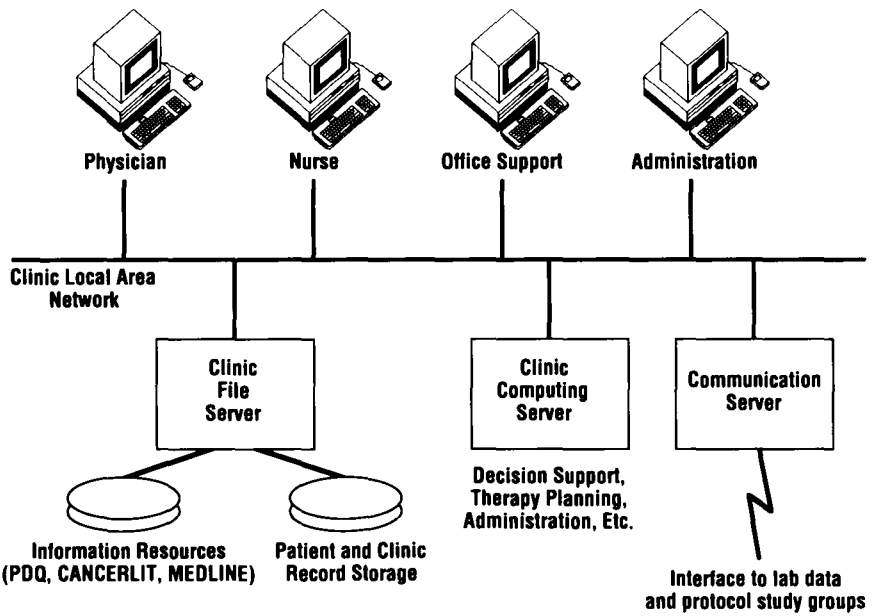

Figure B.2: Distributed Local Server Model

Clearly there is a continuum of system designs between the central model of Figure B.1 and the distributed model of Figure B.2. As suggested in Figure B.2, even in a distributed model, it is possible to share central filing and computing resources where this makes technical and economic sense.

Finally, Figure B.3 illustrates an even broader integration of cooperative services. In order to achieve the most cost effective way to provide clinic services, it may be desirable to have some regional (or even national) centralization of some kinds of services. Through routinely

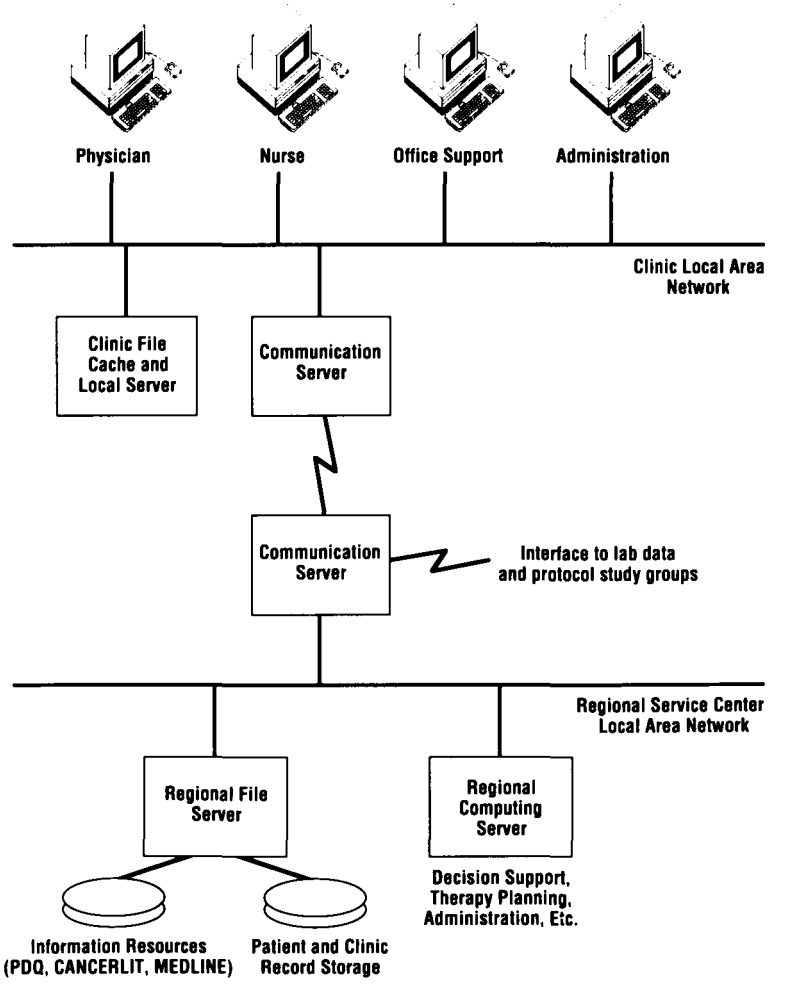

Figure B.3: Distributed Regional Server Model

available communication services ranging from 9600 baud links over dial-up lines to much higher speeds over dedicated leased lines as required, local clinic systems can interface with shared regional systems to provide common file management, decision support computing, or other capabilities. This would be appropriate where it is not possible to provide cost effective computing resources and support for those capabilities in the clinic itself.

As illustrated in the sample distributed system architectures for the oncologist's workstation shown in Figures B.1 - 3, many different configurations of user interface, processing, and storage functions are possible. The most effective one will depend on many details of system implementation, hardware selection constraints, software heterogeneity, and the nature of the target clinical environment. Thus, it is not possible to recommend uniquely a single preferred approach. Rather, the purpose of this discussion is to make clear the range of solutions available with modern computing technology.

NATIONAL CANCER INSTITUTE

International Cancer Information Center
Bethesda, Maryland 20892

Reprinted May 1991