

Computer-aided information management systems in clinical trials

A physician's perspective

Michael N. Pollak *

Department of Medicine, University of Toronto and Princess Margaret Hospital, 500 Sherbourne Street, Toronto, Ontario, M4X 1K9, Canada

Clinical trials involve the administration of new or experimental treatments to patients and the subsequent observation of responses to these treatments over appropriate periods of time. During a clinical trial, large volumes of data describing the course of each patient must be gathered and analyzed. Traditionally, computers have been used only for final statistical calculations after labour-intensive data capture and tabulation. These methods are becoming increasingly expensive, and problems with traditional data management techniques in clinical research are compounded by trends including increasing numbers of patients, increasing length of follow-up period, increasing numbers of relevant treatment and response variables, and participation of geographically dispersed research groups in 'multicentre' trials. Existing computer systems to aid with data management in clinical research are reviewed and criticized, and a new system designed to solve data management problems as perceived by the clinical researcher is described.

1. INTRODUCTION

Sophisticated medical computer applications depending directly on disciplines of Computer Science such as artificial intelligence, mathematical simulation, and information theory have made important contributions to medical practice and research. However, many of the most frustrating information management problems commonly encountered by clinicians are in principle amenable to solution by the application of relatively familiar 'data base management' methodology that is widely employed in industry and commerce. The successful utilization of such methods to manage clinical information has proven to be slower and

more difficult than had been anticipated in the 1970's. This is in part because some systems were designed by computer advocates with idealistic goals such as 'eliminating the paper medical record', rather than being specifically developed to solve in a cost-effective manner the practical information-handling problems perceived by the intended users.

We have designed a computer-based information processing system to solve data management problems encountered in clinical research. 'Clinical research' involves administering new treatments to patients and observing responses over time in order to judge the efficacy of the treatments. Often hundreds of patients on a variety of experimental treatments have to be followed with respect to a large number of 'response variables' for months or years, and statistical comparisons of the various treatment groups are frequently required for ethical as well as scientific reasons.

* M.N. Pollak (MD., F.R.C.P.(C)) is currently a Terry Fox Fellow at the Ontario Cancer Institute, 500 Sherbourne Street, Toronto, Ontario, M4X 1K9, Canada

Surprisingly, there has been relatively little attention given to the application of modern information management techniques to optimize data capture and analysis in clinical research, although it is true that many large clinical research projects use computers for final statistical calculations, after labour-intensive data collection and tabulation steps.

Much clinical research is conducted in the form of 'clinical trials' of new treatments. Typically, the new treatment is described in a 'treatment protocol', which specifies the characteristics of patients that are appropriate candidates for the experimental treatment, the nature of the treatment and how it is to be administered over time, and what clinical and laboratory parameters are to be observed in order to monitor response and/or toxicity. Designs of clinical trials are generally complex: issues related to randomization of patients, control groups, informed consent, and stratification of patients must be specifically addressed. However, the various different clinical trial methodologies generally have similar kinds of information collection and processing requirements, although specific details vary from trial to trial.

The magnitude of effort involved in administering clinical research projects is often underestimated. The National Library of Medicine (USA) has catalogued over 2000 treatment protocols being tested in clinical trials in cancer research alone. Each of these typically involves between 10 and 1000 patients, and each requires collection and analysis of clinical data on every patient involved.

Since formal clinical trials first were employed in the 1940's to study the efficacy of anti-tuberculous drugs [1], they have become generally accepted as the ultimate proving-ground of new medical treatments. While there are continuing controversies with regards to various different ways to organize trials, it would appear that clinical trials of one kind or another will remain the benchmarks for evaluating new medical treatments.

Optimum use of computers to aid with data management in clinical trials is not widespread at present. In this paper, we will outline the reasons for this after reviewing the problems of information management in clinical trials as perceived by the researcher. Some of the successes and failures

of computer efforts in the area to date will be examined, and finally we will describe a new Protocol Management System (PMS) we have designed to assist with data handling in clinical research.

2. INFORMATION MANAGEMENT PROBLEMS IN CLINICAL TRIALS AS PERCEIVED BY RESEARCHERS

2.1. Inadequacy of traditional medical records

As in all scientific research, successful completion of a clinical trial depends in part on careful documentation of procedures performed and results observed. However, unlike laboratory investigation, where this can be done with relatively little distraction, the 'laboratory' of clinical trials is often a busy hospital clinic, where the research component must be closely linked to the medical care function, and where observations are generally recorded somewhat haphazardly in the patients' medical records.

It must be explicitly recognized that under these circumstances, the medical record is being called upon to assume a new role. It is not being used simply to document a single patient's progress as an individual (the traditional function), but also to form the basis for integrating data gathered on many patients in order to study the relative efficacy of different treatments. That is, the data is recorded not simply with the objective of documenting each patient's disease, but also to allow for subsequent evaluation of the treatment given.

While traditional 'free format' paper medical records may or may not represent an adequate and inexpensive way to fulfill the objective of documenting a patient's illness, they clearly do not permit convenient manipulation of accumulated clinical observations to compare groups of patients as is required in clinical research. Thus, a labour-intensive and error-prone step of methodically extracting data from medical records at the conclusion of a clinical trial follow-up period, prior to analysis of the data, has become a common task.

It seems clear that capturing clinical data in a

usable form for subsequent analysis as it becomes available during the course of a clinical trial is preferable to the simple use of the traditional medical record system for clinical research purposes. The use of flowsheets and data collection forms has been helpful, but these measures have not obviated the time-consuming 'data-extraction' step mentioned above.

2.2. Data collection: quality and quantity

There are certain trends in clinical research that tend to exacerbate data-handling problems. The volume of data collected for analysis tends to increase with the length of patient follow-up period and the number of patients followed. Increasingly, clinical trials are being employed to study treatments for the chronic diseases that are becoming major concerns of Medicine. Unlike earlier trials, these studies often must follow large numbers of patients for many years if they are to lead to valid conclusions. For example, in order to compare the efficacy of different adjuvant treatments intended to improve long-term survival of breast cancer patients, it becomes necessary to follow hundreds of patients for many years, particularly if anticipated differences in efficacy are expected to possibly be statistically significant, but nevertheless not large.

Epidemiologists and statisticians face challenges in the *design* of such clinical trials; data processing consultants are now being called on to employ modern data capture and data base management techniques to facilitate their *execution*. This contrasts with the situation a decade ago, where computers were merely used for statistical calculations after data was 'manually' collected, tabulated, and submitted.

Difficulties with collection and tabulation of large quantities of clinical data can lead to important secondary problems for those organizing trials. Often, the additional clerical workload imposed on medical professionals and their assistants in capturing data becomes so time-consuming in the context of a busy practice that accuracy in recording suffers. Furthermore, if information-gathering chores are perceived as too troublesome, organizers of clinical trials find it harder to per-

suade individual physicians to enter patients into research trials. Patients then are treated empirically, rather than being followed for research purposes while they are receiving their medical care.

Other problems are straining traditional data collection methods in clinical research. Many protocols must call for the periodic recording of an unprecedented number of variables for each patient being followed. Studies investigating possible long-term effects of radiation exposure or those designed to identify significant risk factors for developing various diseases exemplify this trend.

The geographical dispersion of 'multicentre' trials complicates data collection as well [2]. One problem involves ensuring that distant participating physicians, who may not have been involved in the design of the trial, are aware of the required observations specified by the protocol. Traditionally, this has been done by mailing formal textual descriptions of the protocol to all participants, and in theory this should be adequate. However, in practice, one observes certain pitfalls. If a physician only has a small number of patients in a given clinical trial, he will not have the protocol details at his fingertips, and he may in fact find the interruption of his busy clinic to consult a protocol manual to be a significant inconvenience. This inconvenience is compounded from the practitioner's point-of-view if he must see, in a realistic example, eight patients in 2 h, two of whom are enrolled in protocol A, one in B, one in C, one in D, and three who are being treated without reference to protocols. The necessity of referring to four manuals within 2 h to ensure that the relevant protocol details are adhered to in each case can be a nuisance, and often either the line-up in the waiting room will increase, or, in the interests of expediency, accurate adherence to the protocol will be sacrificed.

Another, more obvious, problem involving data collection in 'multicentre' trials relates to delays in submission of data from the peripheral sites to the central office for analysis. In theory, this need not be an obstacle, as there is generally no need for the central data base to be updated daily, but in practice co-ordinators of trials often experience significant delays when they have to await the submission of late data.

It should be noted that 'decentralized' clinical trials will probably become more common in the future, because clinical research is increasingly being carried out by university-affiliated practitioners in smaller community hospitals, as well as in traditional academic medical centres.

2.3. Data analysis: flexibility and frequency

Using conventional clinical trial data-handling methodology, interim analysis of results while a trial is in progress is expensive and awkward, particularly if the trial administrator wants to review outcome trends in various different treatment groups and examine these in some detail, indexing by age, sex, or various other variables. Often this is done fully only at the completion of the trial, because of the prohibitive cost of reviewing the data periodically as it accumulates over time.

Certain trials run for needlessly long times because of a lack of availability of interim results. There is now interest in dynamically monitoring clinical trials, to allow trial directors to stop recruitment of new patients as soon as an adequate number of patients has been followed for a sufficient length of time to allow valid conclusions to be drawn [3]. Where applicable, this is a significant way of minimizing the cost of clinical research, but it obviously requires the availability of frequent access to interim results as data accumulates. From an ethical point of view, it is desirable to use interim results to enable the detection of differences between various treatment groups with respect to response or toxicity as soon as possible.

Even though the use of computer systems to process clinical trial data using statistical software is widespread, the data capture and tabulation steps are generally not automated, and the provision of interim results requires repetition of the inefficient step of 'manual' extraction of data from medical records or forms.

Furthermore, the analysis of data must often be delegated to computer personnel who use statistical software to generate reports which are then presented to clinical researchers. This process requires excellent understanding and rapport between those 'asking the questions' and those with actual access to the raw data. Where this rapport

exists, the clinical researchers have relatively effective, if indirect, access to the data. However, it is not rare for the analysis of clinical trial results to be held up by what is perceived by some researchers as the need for 'middlemen' between themselves and the data. Optimum use of currently available data base management systems and query languages could allow statisticians and researchers more direct access to their data, but such systems have rarely been employed in clinical research applications to date.

3. EXISTING COMPUTER SYSTEMS USED IN CLINICAL RESEARCH: ACCOMPLISHMENTS AND PROBLEMS

It is not our intention here to comprehensively review all existing software that is employed to aid with information processing in clinical research, but rather to highlight certain general points regarding existing systems, using a few examples. It must be emphasized that computer-aided data processing enjoys widespread use in clinical trials, and that much has been achieved. However, in most instances there is a gap between 'adequate' data processing support (often achieved) and 'excellent' data processing support, which often seems to be within reach, but rarely has been accomplished to date. Wiederhold [4] has recently reviewed medical data base management systems, including certain applications in clinical research.

3.1. Systems that do too little

Much data processing in clinical research has been done using large computer systems and early versions of commercially available statistical software packages such as SPSS (Statistical Package for Social Sciences) or SAS (Statistical Analysis System). This kind of data processing is successful in accomplishing the limited objectives of performing calculations and summarizing data in a presentable form. However, many of the problems of information processing in clinical trials that frustrate physicians occur at the early steps of data capture and entry, and batch-oriented statistical software is not designed to address these

areas. Furthermore, we have already seen that such systems generally do not make the collected data directly available to the researcher when analysis takes place: often access to the data is by means of a non-interactive program run by computer support personnel.

3.2. Systems that do too much

Large-scale general-purpose computer systems that can handle data entry functions, and maintain an on-line data base that is accessible by an easy-to-use query language are commercially available. However, they often are not appropriate for use in clinical research without modification, not only because of their expense, but also because a carefully customized user interface is required in the clinical trial setting.

Important work in using a comprehensive data base management approach to aid in clinical trials has been carried out. However, at least in one case, the 'Oncocin' system [5], the data base management system forms a subset of a larger package that incorporates artificial intelligence techniques to allow the system to advise the clinician on such matters as formulating a therapeutic regimen, or determining the required interval between patient visits.

In our opinion, the designers of such systems have not had to resort to artificial intelligence techniques to solve user-perceived information handling problems; rather, the designers have provided a 'superset' of functional capabilities that represents an interesting artificial intelligence application, but not a capability which clinicians have identified as being desirable. It may or may not be demonstrated in the future that the use of such advanced systems improves patient care, but at this time many clinicians simply want functional 'no-frill' data management support to assist them in clinical trials administration, without simultaneously becoming involved in separate research addressing the question of the possible usefulness of artificial intelligence aids to patient management.

Some ambitious systems endeavour to enable general practitioners to take over functions previously performed by specialists, or otherwise change the patient load of participating physicians. Such

goals must be regarded as controversial at this time. Systems that are primarily designed to help physicians with the paperwork associated with their usual clinical practice will likely be best received.

3.3. Protocol-specific systems

Many protocol-specific systems have been designed and used successfully on an ad-hoc basis to help with the administration of particular trials. Such systems obviously run into problems in the majority of real-life situations, where groups of patients are being treated according to a variety of different protocols in a single institution. Even where single protocols are used sequentially, users want to learn how to use a system that will be generally applicable, rather than learning how to use a new computer system each time a new treatment protocol is introduced.

3.4. User-oriented systems

Few systems such as the one described in the next section are documented in the literature, and only a few commercially available systems that purport to have the required functions are on the market.

Bill et al. [6] have described a system installed at the Mayo Clinic, and have published a description of its implementation and user acceptance. However, the environment in which this system operates is not typical, in that it depends on the existence of other locally developed software, and on interaction with an 'Institutional Data Base'. While the development costs were estimated at \$35 000, it seems likely that significant additional expenses would be incurred if the system were to be set up elsewhere, independent of the computing resources of the Mayo Clinic.

4. 'PROTOCOL MANAGEMENT SYSTEM' (PMS): A NEW COMPUTER SYSTEM TO AID WITH DATA MANAGEMENT IN CLINICAL RESEARCH

4.1. Design objectives

Our goal was to design a computer system that would employ modern data base management

methods to help solve information handling problems as perceived by those involved in clinical research.

The following summarizes our general design objectives:

1. The system must be cost-effective. It should have all the capabilities required to solve the perceived problems, but the temptation to provide extra features that are 'interesting' but non-essential should be resisted. Artificial intelligence subsystems are specifically excluded, simply because they are not required to provide the functional capabilities desired by the clinical researchers we have interviewed. Indeed, some physicians considered 'computer-aided decision making' as undesirable from their point of view, for reasons related to ethical, medico-legal, and patient-acceptance considerations.

We realize that documenting the possible cost-effectiveness of the system will be difficult. This is because PMS will provide services that were previously absent, rather than simply automating functions that were previously performed by secretarial staff. At the very least, we believe that the impression of initial users with regards to the system's economic justification should be reviewed before the system becomes generally available.

2. The data-entry functions must be easily used by physicians and/or clerical personnel in the milieu of a busy clinic. A screen-based menu format is considered acceptable.
3. The data base access and analysis function should be usable by those physicians who have familiarized themselves with the query language.
4. The system must not be specific to a given protocol: it must be able to simultaneously manage different patients on a variety of protocols. Individual patients must be allowed to be treated on different protocols at different times, and to be followed 'off protocol' when required.

5. While the system is designed primarily to collect and analyse data for the purposes of clinical research, provision must be made for optimum utilization of the collected data to help justify the cost of the system. Sample 'spin-off' benefits include billing, appointment scheduling and other administrative functions not directly related to the clinical research.

6. The system must be applicable to clinical research in both ambulatory patients and in-patients. The capability of serving distant users with remote terminals must be provided.

7. Confidentiality provisions must include password protection, and be equivalent to, or superior to standards of protection present in conventional medical record systems.

8. The software development costs must be minimized by using appropriate commercially available data base management systems. A degree of hardware independence is desirable, and this can also be achieved by the use of such systems.

4.2. Functional capabilities

4.2.1. Accepting and indexing definitions of clinical protocols

The required treatments and observations over time as specified by the protocol are accepted in codified form. Permissible values for a large number of treatment and response variables (typically in the order of 300) are defined. This task is complex, but it would only be done once, at the initiation of a protocol.

At this time, the designer of the protocol also supplies a textual description of it, and the kinds of patients it is intended to treat (so-called 'eligibility criteria'). This is used to enable the system to maintain an index of current protocols, which can later be displayed on request. Participating physicians may use this index to review all protocols active at any given time, to determine if a newly-referred patient is suitable for entry into any ongoing studies.

Responsibility for the approval of newly-proposed protocols should be carefully delegated, and

ongoing clinical trials must be re-evaluated frequently. A computer system that aids in the dissemination of a poorly-conceived protocol would clearly be counterproductive.

4.2.2. Aiding in data collection while patients are being treated

This represents a key function of the system, from the user point of view. While our underlying motivation was simply to provide a facility to conveniently capture clinical data as the trial progresses, we have found that users will make use of the stored data not only for the periodic evaluation of the trial results, but also to aid in the day-to-day management of individual patients. For example, a patient's weight or serum cholesterol may have been entered as required by a protocol evaluating a new drug - this data may be retrieved by a physician at a subsequent follow-up visit to evaluate a patient's progress, long before it is used at the conclusion of the trial to help judge the efficacy of the new drug.

Thus, by entering data into the system, a user is not only contributing to the clinical trial, but also is making data available for easy access in day-to-day patient management. This represents a major 'spin-off' benefit.

Specific functional capabilities related to the routine use of the system in the clinic or office are listed below.

(A) The system must be capable of 'recognition' of patients, given name or number. This implies the ability to display demographic data on a patient, the treatment protocol (if any) currently being used, as well as the number of days elapsed since the patient was entered into the clinical trial. A brief statement of diagnosis and previous treatment also forms part of the identifying information. All this information can be updated or corrected as necessary. The elapsed time variable is automatically kept up to date for each patient by the system.

- (B) The system must be able to provide a 'flowsheet' summary of a patient's course to date. This is generally printed for all expected patients before a clinic session begins, but may be printed for any patient on request.
- (C) The physician may assign a patient to a treatment protocol of his choice, provided that the protocol has been defined for the system. He may also change protocols, or decide to discontinue protocol-defined treatment, but continue to follow the patient. (We anticipate that these capabilities will be used at about 10% of physician-patient encounters, those where major treatment decisions are taken.)
- (D) The system must inform the user of required treatments and observations for a specific patient on a given day, according to the relevant protocol. The objective here is simply to make the physician aware of protocol requirements for a specific patient at a given time. The physician then acts as he sees fit, and if he must violate an aspect of the protocol because of a particular situation, he may do so, documenting his action using the capability described under (F) below. Patient exclusion from the aggregate analysis of the clinical trial, if necessary, is carried out when the trial data is analysed. (Unlike the Oncocin System [5], no effort is made to translate the general rules of treatment as specified in traditional protocol descriptions to specific prescriptions for a certain patient. For example, if a protocol states 'give 10 mg cyclophosphamide/kg patient weight intravenously, provided that the white blood cell count is > 5000 and that there are no contraindications', an Oncocin-like system might review the protocol together with patient data and suggest 'give 700 mg cyclophosphamide'. PMS has more limited objectives - simply to make the protocol statements available to the physician as he would find them if he had to look them up in a manual. Interpretation of the protocol for a specific patient remains the responsibility of the physician, not of the system, because we feel that this is preferred by both the physicians and the patients.)

- (E) The user is allowed to enter values for any observation variables defined for the patient's protocol, even if the recording of these values is not mandatory according to the protocol. Provision is made for the recording of a 'pending' value, to deal with situations where, for example, a blood test has been ordered, but the result is not yet available.
- (F) The user is allowed to enter values for treatment variables to describe treatments actually administered. The system does not assume that protocol-specified treatments are always exactly prescribed, or that these are the only treatments the patient receives.
- (G) On request, the system has the capability to prompt the user to supply missing 'compulsory' data items that have not been entered in the past, in case they have become available since the last visit.
- (H) Physicians may request immediate retrieval of any previously entered data item. As mentioned previously, the ability to answer questions such as 'What was Mr. Smith's last hemoglobin?' represents an important 'spin-off' benefit for many users.
- (I) The system must on request state the maximum period of time that can elapse before the next visit, according to the relevant treatment protocol. In practice, the system responds by reporting a time period such as 'two weeks' or 'one month', and allows the doctor and the patient to agree on a specific date. Extra visits, beyond those 'required' by protocol, can also be booked.
- (J) The system must retain data on 'dropouts', whether due to death or other factors, to allow for their incorporation in subsequent analysis.

4.2.3. Aiding in analysis of collected data

At various interim points, and at the conclusion of a clinical trial, administrators will want to review the aggregate data accumulated. Often, two

or more protocols for treating a single condition will be compared. Note that while many of the capabilities listed in section 4.2.2 are used frequently by doctors and assistants at the time of patient visits, the capabilities related to protocol definition and to data analysis are generally used only by trial co-ordinators, and are used relatively rarely. Nevertheless, it is these capabilities that represent the 'raison d'être' of the system.

Users must be able to use a query language to extract and display subsets of patients of interest, and a report generating capability must be available. Data extracted in this way naturally leads to an analysis of the results of a clinical trial. Provision must be made to enable data to be passed to online statistical packages for further calculations as required.

4.3. Implementation considerations [7]

In order to provide required capabilities, we make use of commercially available software packages that provide data base management and query language functions. As mentioned in the introduction, we feel that the cost-effective solution to many information handling problems perceived by physicians lies in the successful application of currently available hardware and software products. (Many physicians, disappointed with currently used systems, feel that their information-handling problems must await fundamental advances in Computer Science for their solution.)

Systems based on larger microcomputers or minicomputers would be suitable for clinical researchers that are caring for about 500 patients, and are employing about a dozen protocols. We estimate that initial hardware and software costs would be in the \$70 000 (U.S.) range for such an implementation.

A suitable configuration for a dedicated system to serve a patient group of this size would include high-level application-specific software, and a commercially available data base management package that would be compatible with the multi-user time-sharing environment provided by a vendor-supplied operating system. Hardware would include a processing unit, hard disk direct-access memory, a printer, several video terminals

with keyboards, and telephone-network compatible modems. Telecommunication access to a large computer facility would be desirable to gain access not only to magnetic tape facilities for archiving purposes, but also to interactive statistical software that might be required in the further analysis of data collected and presented by PMS.

The system could also be implemented on a larger mainframe, and made available to individual users by means of public data networks. Such an implementation would be more expensive, but would be justifiable if a large number of users were to be serviced.

The cost of a PMS implementation on a dedicated system might well be justified by a group of physicians participating in clinical trials, while a centralized implementation might be appropriate for central organizers of multiple clinical trials, such as drug companies, universities, or government agencies.

In either case, one of the challenges to be faced for the successful establishment of a protocol management system would be the careful integration of the new system with existing 'manual' or computerized medical record-keeping methods. This step must be individualized for each institution, to minimize duplication of data-handling, and maximize user acceptance.

5. SUMMARY

Activity in clinical research will probably continue to increase over the coming years, as newly proposed medical treatments are subjected to clinical testing. We have identified common information handling problems encountered by physicians engaged in clinical research, and designed a computer system to address these problems in a cost-effective manner. The perception of many physicians that data processing support in this area is 'adequate' rather than 'excellent' may be due to the fact that many currently available systems, whether simple or sophisticated, were not specifically designed to solve the problems identified by users. By addressing these problems, we hope that PMS will become a clinically useful tool.

In the past clinical research was carried out without computer-aided data processing. We have reviewed various trends that will tend to increase interest in computer methodology on the part of clinical researchers, and we can expect that further computer systems to aid in clinical trial management will be developed. Such systems may be easier to cost-justify than larger, more expensive 'hospital information systems', because they are specifically targeted to help with some of the most labour-intensive medical information management problems that exist.

Hardware costs are not prohibitive; challenges include keeping software development costs to a minimum, and ensuring that capabilities provided coincide with specifications of intended users. Close liaison between physicians and computer experts is a critical factor in the development of successful systems, and poor results must be expected if those who design programs do not have a thorough appreciation of information-handling problems as perceived by medical personnel. It is the solution of these problems, rather than the mere introduction of new technology, that clinical researchers are seeking.

REFERENCES

- [1] Medical Research Council, Streptomycin in the treatment of pulmonary tuberculosis *Brit. Med. J.* 2 (1948) 1073
- [2] M. Braum et al., Trials and tribulations: thoughts on the organization of multicentre clinical trials *Brit. Med. J.* 81 (1980) 918.
- [3] L. Rubinstein et al., Monitoring clinical trials to terminate accrual *Proc. Annu. Sci. Meet. Soc. Clinical Trials* (San Francisco CA, 1981).
- [4] G. Wiederhold, Data bases for health care (Springer-Verlag, Berlin, New York NY, 1981). [This reference contains a more detailed bibliography listing major medical data base applications that have had use in clinical trials management.]
- [5] E. Shortliffe et al., Oncocin: an expert system for oncology protocol management, *Proc. 7th Int. Conf. Artificial Intelligence* (Vancouver, 1981).
- [6] J. Bill et al., Development of a computerized cancer data management system at the Mayo Clinic, *Int. J. Biomed. Comput.* 9 (1978) 477.
- [7] Further details concerning technical specifications are available from the author.