HELP--A HOSPITAL-WIDE SYSTEM FOR COMPUTER-BASED SUPPORT OF DECISION-MAKING

T. Allan Pryor, Paul D. Clayton, Reed M. Gardner,
Randy Waki, and Homer R. Warner
University of Utah/LDS Hospital
Salt Lake City, Utah

1. COMPUTER IMPLEMENTATION

The HELP system,[1] designed and developed at the LDS Hospital in Salt Lake City, is a data driven hospital and medical information system. Being data driven, every item of data, either medical or administrative, captured by the system results in execution of decision processes pertinent to that particular data item. These processes include diagnostic interpretation of the patient disease state, notification to the appropriate personnel of alarm conditions on a patient, generation of patient specific protocols, distribution of orders, generation of patient reports, capturing of patient charges, updating of patient census information, etc.

HELP is currently implemented on a 4 CPU Tandem computer system. Figure 1 is a block diagram of the configuration. As illustrated in the figure each CPU is configured with 768 K bytes of memory. The current total disc capacity is 288 megabytes. Of these, 128 megabytes are duplicated using mirrored disc volumes. A total of 249 communication ports are attached to the system. Fifty of these ports are allocated to printers; 128 to CRT terminals located on all nursing divisions and in all ancillary departments; and 24 to interface 12 peripheral mini-computer systems, including laboratory computers, patient monitoring computers, multiphasic screening computers, etc. The remaining 38 ports interface developmental terminals located in the Department of Biophysics at the LDS Hospital.

Using the "Non Stop" design of Tandem computers, hardware redundancy and backup is automatically incorporated into the HELP system. That is, no single hardware component failure will cause the HELP system to become inoperable. During normal operation, each CPU functions independently providing minimum response time to the user. If a single CPU were to fail, the load performed by that CPU may be automatically switched to another CPU. In the event of a CPU failure and subsequent backup, while there is increased response time, all functions remain available to the users of the system. During the backup mode the developmental capacity is curtailed to assure minimal response time to the clinical users of the system.

In development of the HELP system design criteria were incorporated into the HELP subsystem modules which take advantage of the Tandem hardware and reduce the dependence of the application programs on the communications hardware interfaced to the system. To develop such criteria the major sources of overhead in the Tandem system were isolated. These overhead sources were during the opening of disc files by the application programs and the creation of new program processes. Since the generation of reentrant code is standard for all Tandem compilers, it was decided to not develop multi-threaded processes. The possible memory savings benefiting from such a design did not compensate for the complexity in development of those program modules. This decision, however, added to the frequency of the process creation function, thus contributing to increased overhead in the system. To overcome this problem a concept of "program pipelining" was developed for use by the HELP system. This concept makes use of the virtual memory and paging capability of the Tandem system. "Program pipelining" implies that frequently used programs are always created in a dormant state by the system ready for attachment to any user. At the creation of that program most of the disc files used by it are opened thus ensuring minimal overhead at the time of attachment of that process to a user. Once a user is attached to that program the system automatically creates in the background a new copy of that program to await a new cycle. Because of the unique nature of the Tandem system this additional copy of the program may be physically located in any of the CPUs, not necessarily the particular CPU to which a calling terminal is interfaced. If, however, there is no immediate request for this program it will ultimately be swapped to disc and reside there in a dormant state waiting for a subsequent request by another user.

In those instances where the program is not created prior to a request, other techniques are used to reduce the overhead experienced by the user. In many instances the program creation and opening of files is performed during input by the user of patient/doctor identification. Thus, as the patient responds to inquiries which specifically identify either himself, the patient, or the doctor for whom a transaction is being requested, the system in the background creates the program requested by the user.

Four major subsystems are available to facilitate the development of application programs on the HELP system. These subsystems include the HELP patient data base management subsystem, the HELP decision-making subsystem, the data entry subsystem and the communications subsystem. Each of these subsystems provides a set of procedures which can be used independently or incorporated through system calls in the development of application programs. In their independent use one may through system defined procedures define data bases, develop and access medical computing logic.

The patient data base subsystem is responsible for: (1) the definition of data items to be stored in the patient data base, and (2) the storage/retrieval of those items. Each data item stored on a patient must be defined to the system using the data dictionary programs of the data base subsystem. In the definition of a data base item three basic structures are available to the user. The first is the definition of fixed field data items for which the user defines precisely the location and extent of every field in the data string stored in the patient's record. The second data structure is a variable length data string where the data within the string represent codes defined by the user in the data dictionary. These codes are combined within the data string using a specific hierarchical syntax. This format is useful for such data as medications where only a small subset of the available drugs are entered in a given data field for a patient. To facilitate the coding of diagnoses the SNOMED (Systematized Nomenclature of Medicine) codes have been translated to this format directly from a tape supplied from the American College of Pathology as part of the data dictionary. HELP decisions are stored as a third data structure. Mixed data structure strings may also be stored. These consist of data in two or more of the basic data structures. The stored data string for all structures has header codes appended indicating the class of data to be stored such as chemistry data, medications, etc., the data structure, the field code within the data class, the length of the string, and the time at which the data was collected. In creating the data dictionary the user also defines any keywords which he desires to be associated to each data item. The format of the data items is also defined; i.e., fixed point, floating point, character string, no value, etc. If the data item is a chargeable item such as a laboratory test, the user defines the condition under which the item is to be charged (on order or on completion), the account number of the item and costs associated with the data item. A 10 byte code is constructed which defines the complete structure of the data item. This code also serves as an index to the files containing the text, keywords, associated HELP modules, etc. This 10 byte code is included in both the HELP module files and the data entry questionnaires which refer to the data item to facilitate the storage or retrieval of data strings containing these data items.

Other basic functions of the data base subsystem are the creation of computerized medical records for patients and the storage and retrieval of information on those patients. These routines give to the user high level constructs through which he may define criteria for access within a given patient record.

Three files are used for the storage of the complete medical record. The first file is a permanent file containing complete demographic information identifying the patient and his financial and insurance status. This file is also used to record on a permanent basis abstract information from the patient's medical record. During an admission to the hospital two additional files are kept - one contains demographic data for the current admission. The second contains all medical information recorded on the patient during his stay in the hospital or an outpatient encounter.

To facilitate retrieval of medical information, the medical data is stored chronologically within each data class. To minimize disc accesses these classes were created to include data commonly used together. Multiple access paths are provided to both of the demographic files. Therefore, searches on patient's name, number, room; doctor number, etc. are quickly performed.

The HELP decision subsystem consists of two components. The first titled HCOM (HELP compiler) is a set of interactive programs through which the user develops the medical logic to be processed by the HELP system. Using either screen prompts or a special command language, the user defines the set of data and logic for a particular decision. The construction of a HELP decision module may be thought of conceptually as the development of a computer program using a problem oriented high level language. The HELP language reflects the need to perform sophisticated retrievals from the patient's data base. The criteria which one may include in defining a patient data search include conditional criteria based on other data items, timed events from the patient's data base, Boolean time constraints, or combinations of each. The requested search can be constructed to return not only a single item but multiple items or a derived value. These derived values can either be system defined, e.g., maximum, minimum, first, last, median, mode, etc., or specified by user defined logic written as part of the data search. Within the logic structure of the decision module, the user develops decision criteria based on Boolean constraints, arithmetic and/or statistical models, or causes the HELP system to interact with user defined processes. Each decision module may have from one to ten results. These results may be distributed to multiple locations including back to the calling program, the patient file, or one or more special alert files.

The HELP interpreter is the process which performs the interpretation of all HELP decision modules and transmits the results to the locations specified by the decision module and the calling program. Activation of HELP decision modules is generally through user constructed lists linked to the storage of patient data. Because of the high degree of interaction between the HELP interpreter, the patient data storage and the patient data retrieval processes, these three processes are resident in the system. Therefore, in the storage of a data item which is linked to a set of HELP decisions, no additional process creation is required to accomplish the coordinated tasks of HELP interpretation, data storage, and data retrieval.

Since the decision modules themselves are records in a file, security is easily provided which ensures that only the author of a decision block has the ability to modify decision criteria within his own block. The structure of a HELP module has been designed to ensure a minimal number of accesses to either the patient or data dictionary files for proper interpretation of the HELP logic. For example, included in the HELP module are the complete 10 byte coded definition of the structure and location of data items to be searched in the patient data base.

To decrease the time to process a HELP block two special preprocessing steps are applied to the block as it is compiled for on-line use. They are: (1) a recording of all of the HELP modules which ensure that all information required by a HELP module is available through either a data search or the result of a HELP module which had been processed prior to the processing of this module, and (2) the scanning of all the HELP modules within a block to determine the types of data and time constraints on those data elements which will be required for processing of that HELP block. At time of execution of the block the HELP interpreter uses this information to formulate a single patient data retrieval request particular to the time constraints and data of the HELP modules in that HELP block. The individual data searches requested by a HELP module are then performed from the in core buffer retrieved at the start of the block processing without the need for multiple time consuming disc retrievals.

The data entry subsystem provides a program independent method for development of medical or administrative manual data entry screens. Through an interactive program a user may create a questionnaire consisting of one or more screens. The screens are sequenced by the fixed order of creation or a calculated order based on Boolean logic contained within the individual screens. There are several types of basic screens provided for creation of a questionnaire. They are: (1) a multiple choice screen where the multiple responses are defined, (2) a numeric screen through which one or more numeric values are entered at the time of execution, (3) a free text screen for entering of ASCII characters into the patient's record, and (4) time screens constructed to enter a time value into the patient's record. Additional pseudo screens may also be used. These pseudo screens are of the same type previously mentioned but are not shown to the user at time of execution of the questionnaire. They have the facility to construct automatically the response to that question and store that response within the patient's data record without being presented to the user. An example of the use of a pseudo screen might be the inserting of the time of data entry. If the user chooses to develop the questionnaire where the time recorded as the data entry time is always the current time the pseudo screen could be used to enter this information without the need of requesting the time from the user. In all instances of screen types, diagnostics are available to prevent input of data which is inconsistent with the type of screen defined and the values allowed for numeric entry. At execution instead of viewing every screen, an input command string may be used. In this case the user on responding to the first screen types a command string. Upon detection of the command string the system uses the data in the string as the entries from the individual screens. If errors are detected the screen where an error was detected is displayed and the user continues normally. An example of the use of a command string could be in the ordering of a SMAC blood test at the next routine collection of blood samples. In this instance, rather than be shown the screens presenting the list of possible blood tests and times associated with those tests, the user would enter a command string giving responses to the screens requesting the laboratory test to be ordered, the type of test and the desired collection time. On termination of that command string the user is presented with a review of the responses for verification. Upon verification the data string is stored in the patient's record.

The terminal handler subsystem makes possible the development of application programs which are independent of the terminals from which they are activated. This subsystem resides between the user's terminal and the application program to which he is attached. It performs several special functions for the user. The first is to determine the program availability to a specific terminal. This is done through a series of files defining the characteristics of each terminal. Among the characteristics defined is the menu screen to be displayed on that terminal. This menu screen gives the set of options available at that terminal and ensures that no unauthorized program or data may be accessed from that terminal. The terminal file also contains information on the physical characteristics of the terminal. The physical location of the terminal and the location of printers closest to the terminal are also stored in the terminal file. The location information is used in routing reports to the proper location.

Another function of the terminal handler is to capture information from the user before attaching that user to the application program requested. The three basic items of information are the patient identification, the doctor associated with that patient, and the user who is currently performing the program request. Not all of these items of information need, however, be captured at any program request. In the definition of the menu screen the developer defines the application programs to be presented on the menu screen. For each program presented on the screen the developer also defines what front end data is to be captured. The creation of the requested process is performed in the background parallel with the acquisition of the front end data. In addition to the patient ID, doctor, user ID, the HELP data base management process to be used by the application program is also transmitted to the application program as part of the creation function. Using the HELP system a number of applications have been developed and are in use at the LDS Hospital. Table I is a list of the major applications currently in use at the LDS Hospital. In virtually all of these applications some level of medical decision-making has been incorporated to allow the computer to serve not only as a tool for recording and transmission of information, but the provider of new information to aid in the care and diagnosis of the patient. The HELP system, while still undergoing continued improvement, has proven the utility of computerized medical decision-making to solve some of the complex problems of medical care.

2. DECISION ANALYSIS USING PATIENT SPECIFIC EXPECTED UTILITIES

It is the purpose of this section of the paper to describe how computerized decision analysis is used to assist the physician in treating patients with suspected ischemic heart disease. Unique features of the HELP system allow decision analysis to be implemented in a simple, useful manner.

Decision analysis is one of several quantitative approaches to the decision-making process. When this method is used, the expected utilities for alternative courses of action (or inaction) are calculated, and the choice with maximum expected utility is selected. To calculate expected utilities, a decision tree is constructed with branches which lead to all outcomes which might result from alternative choices of the decision-maker. Each branch point in the tree is referred to as a node. Nodes whose branches represent alternative choices to the decision-maker are decision nodes. Those nodes whose branches represent uncertain events beyond the control of the decision-maker (test results, response to therapy, etc.) are called probability or chance nodes. In the clinical setting the branches of the decision tree typically lead to a variety of outcomes. These outcomes range from a state of normal health obtained with no (or inexpensive and riskless) therapy to the worst case of death following a painful illness and costly therapeutic measures. After the relative utilities (values) of each of the alternative outcomes are established, the probabilities that a specific patient would arrive at each outcome must be calculated. Next, the products of the probability that a specific outcome will occur times the value of that outcome are summed for all outcomes that branch from one choice at a decision node. This sum is the expected utility associated with making that choice.

Decision analysis is an attractive approach to decision-making because probabilities that outcomes will occur and judgments of the values of alternative outcome are stated explicitly and separately. A choice based on maximum expected utility may disagree with an expert's opinion, if that expert is inconsistent, the probabilities in the decision tree are inaccurate, or the utility assignments are not appropriate. It is possible to isolate which stages of a decision process are most crucial by using sensitivity analysis. To accomplish this parameters or utilities are varied over an acceptable range of uncertainty, and the threshold values at which one decision alternative becomes preferred over others are determined. The use of decision theory(2-8) has been applied to actual or simulated case histories in patients with coronary artery disease,(9,10) acute renal failure,(11,12) pulmonary emboli, (13) appendicitis, (14,15) space occupying lesions on intravenous pylorograms (16) and a variety of other problems. The actual use of computerized decision theory in clinical practice has been accomplished in the treatment of Hodgkins disease.(17)

2.1 METHODS

The HELP system for medical decision-making(14) is well-suited for implementation of decision analysis for two reasons. The first reason is that medical logic for each separate node or outcome in a decision tree can be constructed modularly using medical language. The tree can be expanded or pruned simply. This capability allows medical or technical personnel with minimal training to modify the logic. The second advantage of the HELP system stems from the fact that it is interfaced to a comprehensive clinical data base. This means that the decision criteria can be tested using real patient data prior to implementation in a clinical setting. The clinical information necessary to evaluate a proposed decision tree (with the exception of preference data needed for utility estimates) is presently being entered into the computer as part of routine clinical service. The resultant data base consists of computerized ECG interpretations, hemodynamic measurements, history of symptoms, location and severity of coronary artery lesions, results of stress tests, follow-up concerning quality and length of post-surgical life. Figure 2 is a partial representation of the decision tree which has been implemented with the HELP system. The branches which are not shown are structurally similar to the example except that costs and probabilities for analogous outcome status reflect the fact that a stress test was or was not performed, etc. The entire decision tree contains 51 unique HELP sectors which serve to evaluate the utilities of given outcome states. They are denoted by U's and are at the tips of the tree branches. There are 25 chance sectors (denoted by C's) which determine probabilities that a particular test result or outcome state will occur. There are two branches from each chance node; the probability (p) of the upper branch is calculated according to the knowledge and rules contained in the chance HELP sector for the node. The probability of the lower branch is then 1-p. The expected utility of each branch is calculated as the product of the probability that that outcome will occur times the value for the utility of that outcome.

There are 25 Decision nodes for which HELP sectors have already been written. The logic in the Decision node sectors calculates the sums of the expected utilities for all outcomes of the two alternative actions. The action for which the value of the expected utility is largest is the optimum procedure. Each decision sector is activated when data referenced in the sector is added to the patient record or is called by some sector to the left of it in the tree. Thus, this tree will be activated whenever (a) a history of typical angina, (b) the results of a stress test, or (c) the results of a selective coronary arteriogram are entered into the computer. Only the relevant portion of the tree will be activated whenever data entry occurs. This activation always occurs at an input to a decision node; before the decision sector can be evaluated it must call the probability and utility sectors. Whenever additional knowledge (in the form of test results) is obtained the sector logic is re-evaluated. Thus, decisions which are optimal when the only thing known about the patient is that he has typical angina, may not be the optimal decisions after the results of selective coronary arteriograms are known.

This tree illustrates how it is possible to construct the decision analysis tree in a modular fashion. One HELP sector exists for each chance, decision or utility node. Using HELP, the logic criteria and medical knowledge in a single sector can be updated and modified independently of the remainder of the tree, using natural language for medical terminology. The purpose of these sectors is to combine the knowledge of local experts, results from medical literature, and conditional probabilities obtained from the patient data base. This patient specific logic will enable the chance sectors to predict the likelihood of each outcome and assure that the utility sectors represent the values of each state to the particular patient being analyzed.

2.1.1 Typical Utility Sector - Sector 3 (see figure 3)

In all the explanations of HELP sectors which follow, the word "SEARCH" appears in two contexts: (1) Immediately following a logic item (e.g., item C). This format directs the HELP interpreter to search the patient record for the desired data. A search is performed every time "SEARCH" is encountered in this context; (2) Imbedded in the item (e.g., item F). This format tells the HELP interpreter that the data can be found in a special search sector that searches the patient record for data needed by several sectors before those sectors are evaluated. This approach is useful for items such as age and sex since the HELP interpreter will only search for the data once and then will make these results available to other sectors requiring them.

The utility sector 3 (figure 3) computes the expected utility of a final outcome in which a patient with typical angina, a positive stress test and triple vessel disease by selective coronary arteriography (SCA) receives surgical treatment. Nearly all utility sectors have this general form.

Item A (0.8) is an estimate of the average quality of life following triple graft surgery. This number was obtained from analysis of responses to a questionnaire sent to patients who had undergone this operation. Using this scale, a value of 0.0 represents death and 1.0 represents a morbidity free state.

Item B selects the appropriate annual coronary mortality rate based on whether or not the patient has left ventricular contraction abnormalities. If item M exists (indicating normal contractions), the mortality rate is 3.73%. If search item N exists (indicating abnormal contractions), the mortality rate is 2.86%. Otherwise the mortality rate is 3.27% (nature of the contractions is unknown). The mortality rates are from Hall 1980. (18)

Item C retrieves the value of normal life expectancy which was calculated in another sector. This value is based strictly on the patient's age and sex.

Items D and E compute the patient's adjusted life expectancy. An exponential survival curve with the mortality from item B is integrated across the remaining life expectancy of a normal person with the patient's age and sex according to the following equation:

$$\frac{1 - \exp(-\text{coronary mortality rate} * \text{normal life expectancy})}{\text{coronary mortality rate}}$$

Item F calculates the total expected benefit in dollars of being in this state. Adjusted life expectancy (item E) is multiplied by the quality of life (item A) and by the estimated dollar worth of one quality year of life (search item C, which is $25,000).

Item G estimates the total expected cost in dollars of being in this state. The cost of stress testing (search item G which is $130) is added to the cost of the SCA (search item H which is $1500) and the cost of surgery (search item I which is $15,000).

Final Evaluations. Final evaluation B is calculated as the expected utility for the patient by being in this state. The costs (item G) are subtracted from the benefits (item F) and then multiplied by a constant (search item D which is 1). Final evaluations A and C are used in sensitivity analysis. A multiplies quality of life by 1.25 (search item A) while C multiplies qualify of life by 0.75 (search item B) in order to display upper and lower extremes of the utilities generated.

2.1.2 Typical Chance Sector - Sector 120 (see figure 4)

The chance sector 120 computes the probability that a patient with typical angina and a no stress test will have coronary artery disease (CAD) determined by selective coronary arteriography (SCA). Each item of logic is explained below.

Item A. This item causes the HELP interpreter to jump to item E if the patient is male (search item L is zero for females and 1 for males).

Item B. This item selects one of four probabilities depending upon the age of the female patient (search item K).

Item C. The probability from item B is assigned to item E for use in the final evaluation.

Item D. This item tells the HELP interpreter to skip item E and go directly to the final evaluation.

Item E. This item selects one of four probabilities depending upon the age of the male patient (search item K). The probabilities in this sector come from Diamond 1979.[19]

Final Evaluation. The value of item E is now the patient-specific probability of CAD given typical angina and a no stress test based on age and sex. The HELP system stores this as the final evaluation of sector 120 for possible use by the other sectors.

2.1.3 Typical Decision Sector - Sector 202 (see figure 5)

The decision sector 202 decides whether or not to perform a SCA on a patient with typical angina and a positive stress test. The expected utility of node 202 will be returned as the final evaluation for use by sector 201 (see the FINAL EVALUATION). If activated by a positive stress test, display of the message will depend on which decision has the highest expected utility. Each item is explained below:

Items A-N. These items access the final evaluations of the utility, chance and decision sectors required by this sector which have already been executed.

Item O. This is the expected utility of performing a SCA on a patient with typical angina and a positive stress test. The expected utilities given by sectors 203-208 (items H-M) are multiplied by the probability of each type of CAD given by sectors 103-107 (items C-G). This gives the expected utility of performing a SCA assuming the patient has CAD. The expected utility of performing a SCA assuming no CAD is given by sector 15 (item A). These two expected utilities are multiplied by the probabilities of CAD and no CAD, respectively, given by sector 102 (item B and 1-B).

Final Evaluations. Final evaluation A of this sector is the expected utility of node 202. This is the maximum of the expected utilities of performing a SCA (item O) and not performing a SCA (item N). Final evaluation B is the difference between the utilities of the two decisions (item O minus item N).

Feedback to the user regarding the optimum decision is constructed so that the difference of the utilities is given; e.g., "the expected utility of performing a stress test is $1500 greater than the expected utility of no stress test." This information gives the user a quantitative feeling of how strongly the utility of the best decision dominates the alternative choice.

By using multiple final evaluations in the HELP sector logic, sensitivity analysis is automatically performed for the decisions regarding surgical vs medical treatment. Our present estimates of utility are patient specific only to the extent that they are based on the age, sex and severity of disease. Both quality-of-life and the dollar value of one year of morbidity-free life require subjective estimation.[4] Our present estimates of values for the quality-of-life were obtained from responses to follow-up questionnaires returned by the study subjects. Sensitivity analysis is performed by computing nine possible values of the decision sector. In our initial formulation of the expected utility sector we calculate the dollar value of one year of life as the product of the dollar value of one year of morbidity-free life times the expected quality-of-life which is expressed as a decimal fraction between 0 (dead) and 1 (no morbidity). We obtain three values of the expected utility for a state (maximum, mean, minimum) by multiplying the mean expected value by 1.25, 1.0, and 0.75 respectively. There are nine resulting combinations for comparison between the surgical and medical treatments. The decision sector message is structured to indicate for which of the nine cases surgical treatment would be expected to yield maximum utility and for which of the cases medical treatment is preferred.

We have run the decision analysis tree on the 729 patients for whom the required data (including physician and hospital costs and salary before and after surgery) is present in the computer. All of these patients did undergo surgery but it is conceivable that alternative therapy could have resulted had they seen a different physician. Therefore while running the analysis on this group of patients, histograms have been constructed for each decision node. These histograms show the difference in expected costs between alternative one and alternative two along the horizontal axis and the frequency of patients for which the difference was calculated on the vertical axis. On this histogram a value of zero shows that the alternatives are equivalent in terms of expected utility. Values less than zero show that alternative two is favored.

2.2 RESULTS

The modular construction of a decision tree using HELP sectors has proved to be a feasible approach. Figures 6a and 6b show histograms of the differences in expected utility between surgical and medical treatments for 128 patients with single vessel disease in the left anterior descending coronary artery. Figure 6a assumes the maximum estimate for quality of life following surgery and the minimum estimate of qualify of life for medical treatment. In figure 6b just the opposite is assumed. The recommended treatment for every patient in this group is changed by varying the estimates of quality of life. Similar results are obtained for single vessel disease at other locations. However, in the instance of double, triple or left main disease, the recommended treatments are not so sensitive to estimates of quality of life and not all these patients change categories even when the extreme circumstances are evaluated.

2.3 DISCUSSION

The decision tree for patients with typical angina is similar to those developed for the same purpose by other investigators.[9,10] However, the use of HELP to construct such a tree is a flexible approach which allows a variety of trees to be implemented using the modular fashion for the decision logic as presented in this paper. At this stage of our investigation the decision logic is rudimentary and is being updated to reflect new information on the subject as it appears in the literature.

The analysis appears to be sensitive to the appropriate parameters. The decisions are not sensitive to the costs of tests and therapies unless the outcome, quality of life, or survival probability are influenced by them. In the case of single vessel disease, mortality rates for surgical and medical treatment have not been shown by other researchers to be different. Therefore, the expected quality of life following treatment is the predominant parameter in the decision process. However, in multiple vessel or left main disease, published reports have shown a difference in mortality with surgical vs medical treatment. This difference was expressed in our logic. In such cases, quality of life variation did not affect the treatment of choice for all patients. For younger patients who would have a longer expected lifespan, different quality of life levels resulted in larger differences in utilities between alternative treatments.

In summary, we have presented this work to show that sophisticated approaches to the complex process of medical decision-making can be performed using HELP. Much of the hard data needed for analytical methods is not immediately available, but decision analysis can identify which data are most needed and provide a means so that medical knowledge will be applied uniformly to all patients serviced by the system.

REFERENCES

1. Warner HR, Olmsted CM, Rutherford BD. HELP--a Program for Medical Decision-Making. Comp and Biomed Res 5, 65-74 (1972).
2. Raiffa H. Decision Analysis. Reading Mass., Addison-Wesley (1968).
3. Kassirer JP. The Principles of Clinical Decision-Making: An Introduction to Decision Analysis. Yale J Biol Med 48:149-164 (1976).
4. Albert DA. Decision Theory in Medicine. A Review and Critique. Milbank Mem Fund Q 56: 362-401 (1978).
5. McNeil BJ, Keeler E, Adelstein SJ. Primer on Certain Elements of Medical Decision-Making. N Engl J Med 293:211-215 (1975).
6. Patton DD. Introduction to Clinical Decision-Making. Semin Nucl Med 8:273-282 (1978).
7. Ransohoff DF, Feinstein AR. Problems of Spectrum and Bias in Evaluation of the Efficacy of Diagnostic Tests. N Engl J Med 299:1259-1263 (1976).
8. Schwartz WB, Gorry GA, Kassirer JP, Essig A. Decision Analysis and Clinical Judgment. Am J Med 55:459-472 (1973).
9. Pauker SG. Coronary Artery Surgery: The Use of Decision Analysis. Ann Intern Med 85:8-18 (1976).
10. Weinstein MC, Pliskin JS, Stason WB. Coronary Artery Bypass Surgery: Decision and Policy Analysis. Costs, Risks, and Benefits of Surgery, Oxford University Press (1977).
11. Gorry GA, Kassirer JP, Essig A. Decision Analysis as the Basis for Computer-Aided Management of Acute Renal Failure. Am J Med 55, p 473 (1973).
12. Pliskin JS, Beck CH. Decision Analysis in Individual Clinical Decision-Making: A Real-World Application in Treatment of Renal Disease. Meth Inform Med Vol. 15, pp 43-46 (1976).
13. Pauker SG, Kassirer JP. Clinical Application of Decision Analysis: A Detailed Illustration. Semin Nucl Med 8:324-335 (1978).
14. Warner HR. Computer-Assisted Medical Decision-Making, Academic Press, New York (1979).
15. Neutra R. Indications for Surgical Treatment of Suspected Acute Appendicitis in Costs, Risks and Benefits of Surgery. Bunker, Barnes, Mosteller eds. pp 277-307, Oxford University Press, New York (1977).
16. Fryback DG, Thornburg JR. Informal Use of Decision Theory to Improve Radiological Patient Management. Radiology 129:385-388 (1978).
17. Safran C, Desforges JF, Tsichlis PN, Bluming AZ. Decision Analysis Applied to Lymphangiography in Hodgkin's Disease. N Engl J Med 296:1088-1092 (1977).
18. Hall RJ, Mathur VS, Garcia E, DeCastro CM. The prolongation of Life by Coronary Bypass Surgery. In Hurst JW (ed), "Update II: The Heart", McGraw-Hill, pp 175-204 (1980).
19. Diamond GA, Forrester JS. Analysis of Probability as an Aid in the Clinical Diagnosis of Coronary Artery Disease. N Engl J Med 300:1350-1358 (1979).

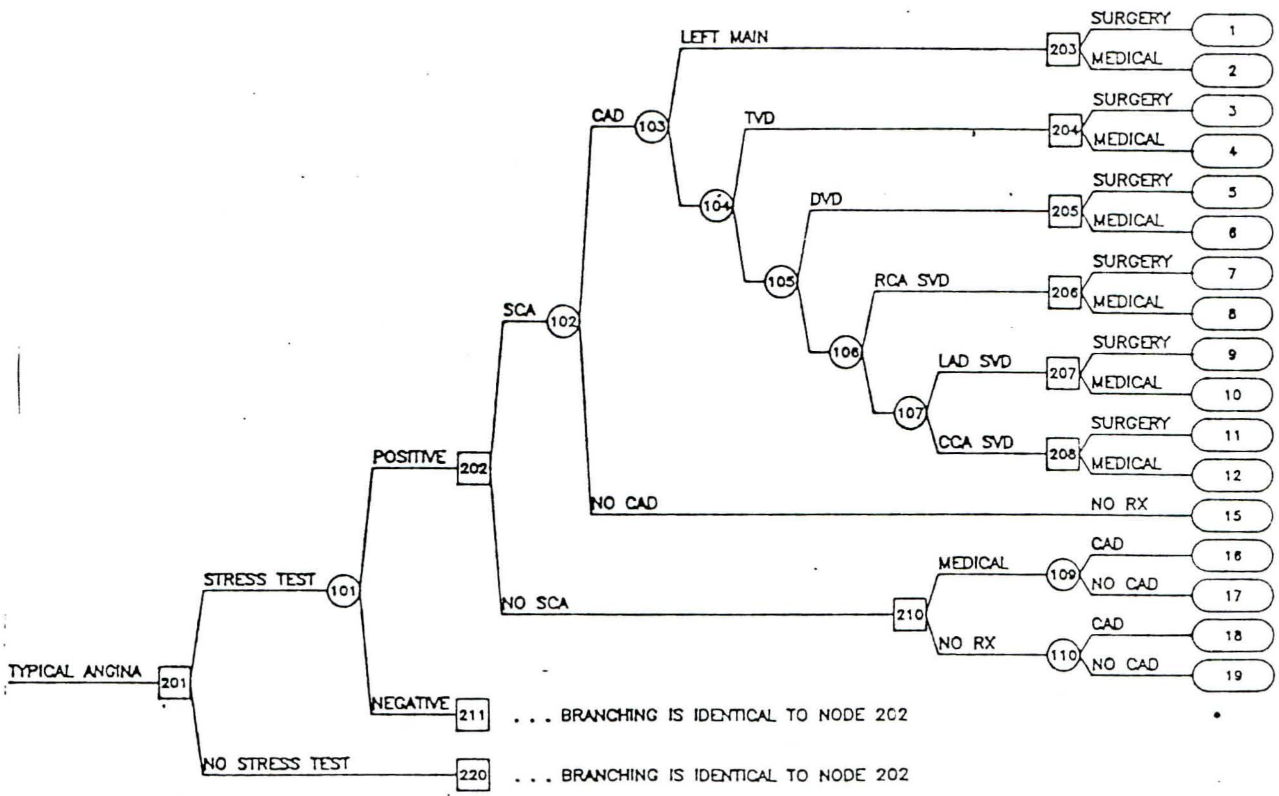

Figure 2 The decision tree for typical angina. Square nodes represent decisions, circular nodes represent chance events (patient specific probabilities) and oval nodes represent outcomes.

SECTOR 3 DECISION TREE FOR CORONARY ARTERY DISEASE
Utility : Typical angina, + stress, TVD by SCA, surgery

FINAL EVALUATIONS:
A VAL: SEARCH VAL D *(SEARCH VAL A * F - G)
B VAL: SEARCH VAL D *(F - G)
C VAL: SEARCH VAL D *(SEARCH VAL B * F - G)

SECTOR LOGIC:
A ARITH: 0.8
B ARITH: IF SEARCH EX M THEN 0.0373 ELSE(IF SEARCH EX N THEN 0.0286 ELSE 0.0327)
C SEARCH: (A) (251) Normal Life Expectancy, USE: -(ITEM VAL B * A / 4)
D ARITH: EXP(-B * C)
E ARITH: (1 - D) / B
F ARITH: E * A * SEARCH VAL C
G ARITH: SEARCH VAL G + SEARCH VAL H + SEARCH VAL I

Figure 3 Utility sector 3. See text for a detailed explanation

```
ECTOR 120      DECISION TREE FOR CORONARY ARTERY DISEASE
               (1)  P( CAD by SCA : Typical angina, no stress)


          FINAL EVALUATIONS:
          A   VAL:   E

          SECTOR LOGIC:
          A   ARITH:     IF SEARCH VAL L EQ 1 THEN GO TO E
          B   ARITH:     IF SEARCH VAL K LT 40 THEN 0.258 ELSE (IF SEARCH VAL K LT 50 THEN 0.55
                         ELSE(IF SEARCH VAL K LT 60 THEN 0.794 ELSE 0.906) )
          C   ARITH:     E = B
          D   ARITH:     GO TO FE
          E   ARITH:     IF SEARCH VAL K LT 40 THEN 0.697 ELSE(IF SEARCH VAL K LT 50 THEN 0.873
                         ELSE(IF SEARCH VAL K LT 60 THEN 0.92 ELSE 0.943) )
```

Figure 4 Chance Sector 120. See text for a detailed explanation.

```
:CTOR 202      DECISION TREE FOR CORONARY ARTERY DISEASE
               Expected utility : Typical angina, + stress

          FINAL EVALUATIONS:
          A  VAL:   MAX(O, N)
          B  VAL:   O - N

          SECTOR LOGIC:
          A  SEARCH:     (A)  (15)  Utility : Typical angina, + stress, no CAD by SCA, no treatm
                         (FE B)
          B  SEARCH:     (A)  (102) P(  CAD by SCA : Typical angina, + stress)
          C  SEARCH:     (A)  (103) P(  LM CAD by SCA : Typical angina, + stress, CAD by SCA )
          D  SEARCH:     (A)  (104) P(  TVD by SCA : Typical angina, + stress, non-LM CAD by SCA
          E  SEARCH:     (A)  (105) P(  DVD by SCA : Typical angina, + stress, non-LM non-TVD
                         CAD by SCA )
          F  SEARCH:     (A)  (106) P(  RCA SVD by SCA : Typical angina, + stress, SVD or no
                         significant lesions by SCA )
          G  SEARCH:     (A)  (107) P(  LAD SVD by SCA : Typical angina, + stress, LAD SVD or
                         circumflex SVD or no significant lesions by SCA)
          H  SEARCH:     (A)  (203) Expected utility : Typical angina, + stress, LM CAD by SCA
          I  SEARCH:     (A)  (204) Expected utility : Typical angina, + stress, TVD by SCA
          J  SEARCH:     (A)  (205) Expected utility : Typical angina, + stress, DVD by SCA
          K  SEARCH:     (A)  (206) Expected Utility : Typical angina, + stress, RCA SVD by SCA
          L  SEARCH:     (A)  (207) Expected utility : Typical angina, + stress, LAD SVD by SCA
          M  SEARCH:     (A)  (208) Expected utility : Typical angina, + stress, circumflex
                         SVD by SCA
          N  SEARCH:     (A)  (210) Expected utility : Typical angina, + stress, no SCA
          O  ARITH:      B *(C * H +(1 - C) *(D * I +(1 - D) *(E * J +(1 - E) *(F * K +(1 - F)
                         *(G * L +(1 - G) * C) ) ) ) ) +(1 - B) * A
```

Figure 5 Decision sector 202. See text for a detailed explanation.

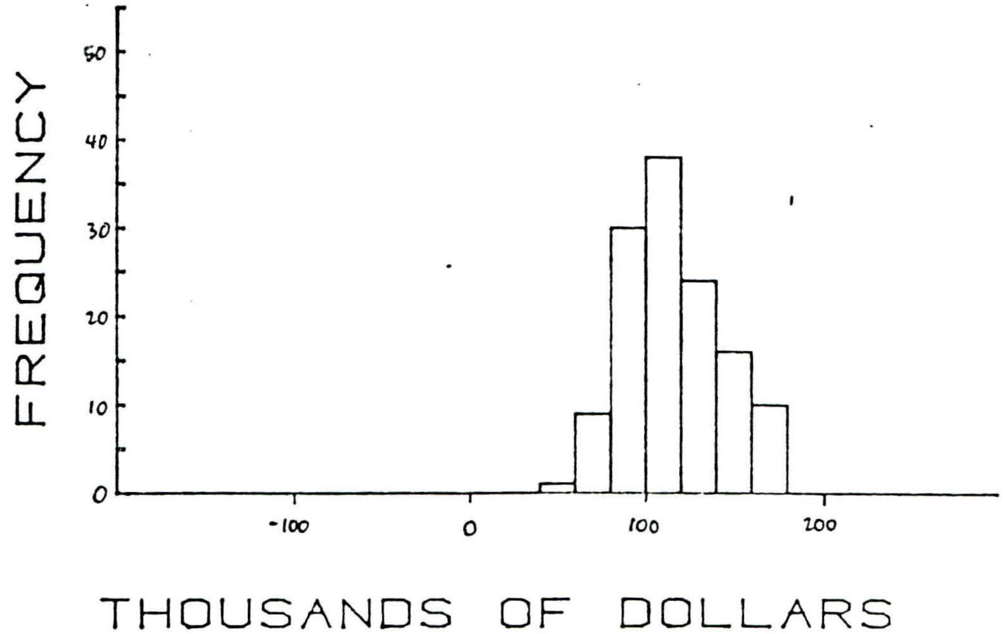

Figure 6a

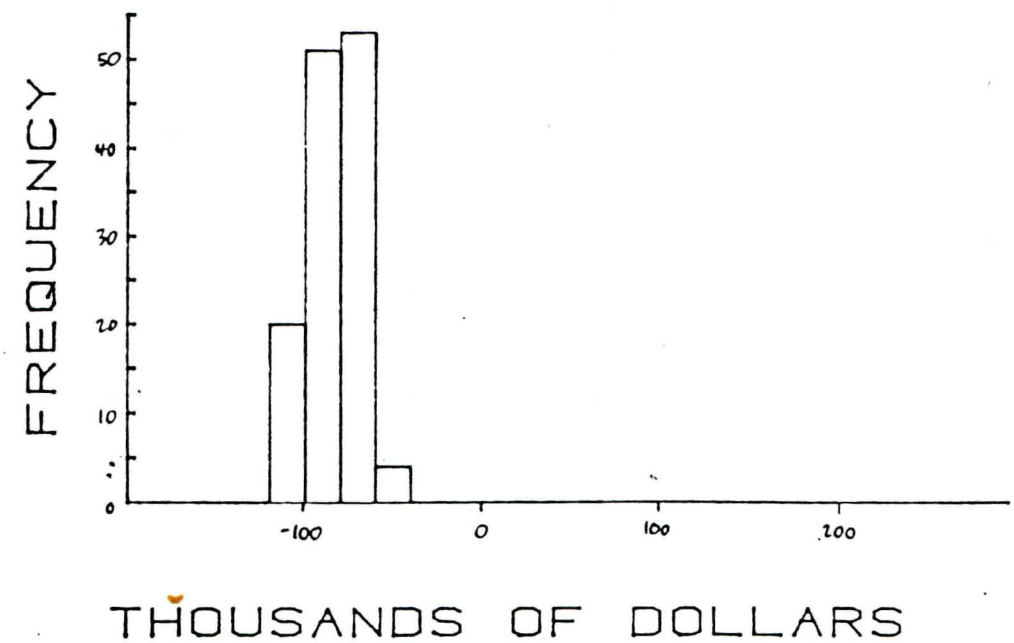

Figure 6b

Fig. 6. Sensitivity analysis of quality of life (QOL) for the decision to treat surgically or medically 128 patients with LAD SVD (no stress test). The horizontal axis denotes the difference between the expected utilities of the two treatments ($20,000 per bin). Fig. 6a shows the effects of decreasing post-surgical QOL by 25% and increasing QOL under medical treatment by 25%. Fig. 6b shows the effects of the reverse change.

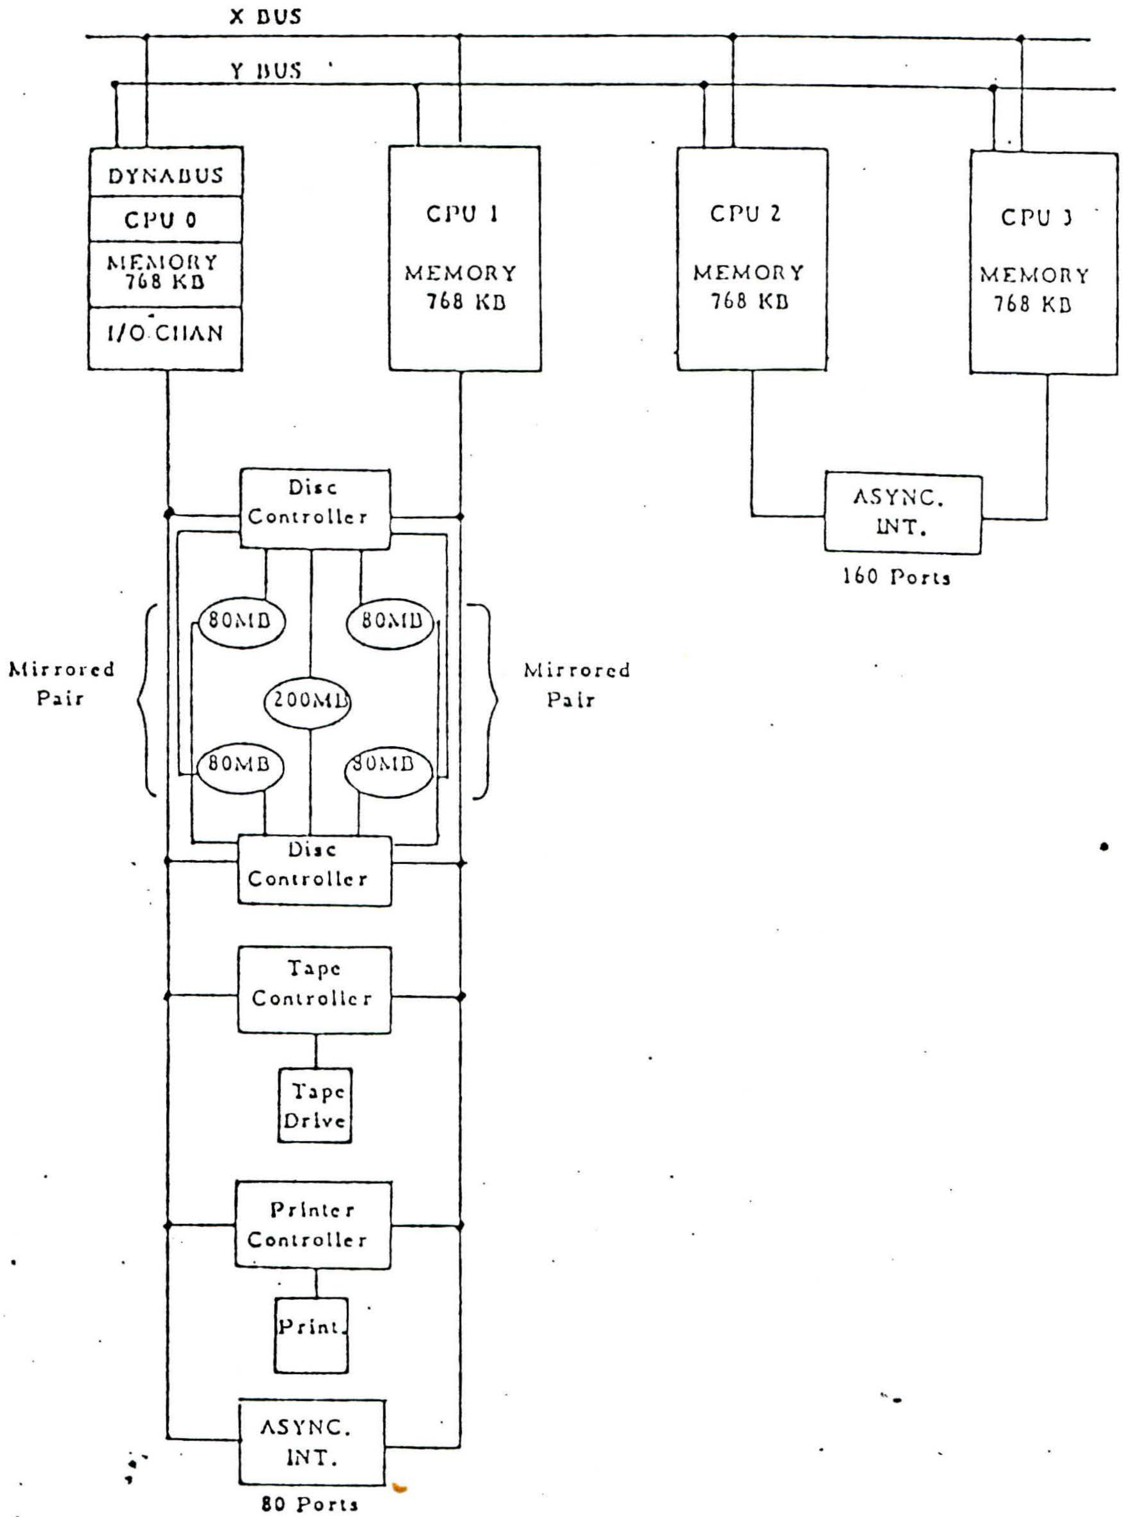

Figure 1 Block Diagram of Tandem Computer System for HELP

13000 RAYMER STREET
875-0555 • 933 1377

APPLICATIONS

- o ADT
- o Order Entry/Charge Capture
- o Pharmacy
- o Clinical Lab
- o Radiology
- o Blood Gas
- o Pulmonary Functions
- o Heart Catheterization
- o Cardiology/EKG
- o Hemodynamic/ECG Monitoring
- o Medical Records
- o Peer Review
- o Research

Table I List of Applications on the HELP System