History of the Development of Medical Information Systems at the Laboratory of Computer Science at Massachusetts General Hospital

By G. Octo Barnett, M.D.

Professor of Medicine, Harvard Medical School
Physician, Director of Laboratory of Computer Science
Massachusetts General Hospital

The reconstruction of history is always a difficult task. The important decisions, the critical issues never seem to be adequately documented and cannot easily be recalled. It is difficult enough in affairs of national importance to try to establish "who knew what, and when did he know it?" In considering the history of the activities of the Laboratory of Computer Science at Massachusetts General Hospital, I have found it equally difficult to try to establish: "what did we do, and why did we do it?"

I have little confidence in my judgment as to what were the important contributions of the Laboratory of Computer Science in the first decade of its existence, and even less confidence in my ability to recall and to interpret appropriately the contributions of a host of individuals who participated in the effort.

The only option seems to be to frame this presentation as a series of flashbacks consisting mostly of hazy memories of incidents and selected extracts from grant requests and progress reports. I do not pretend that these flash-backs are either a valid scientific quest or a valid historical review, and I fully appreciate that these memories are greatly colored by the events and the activities that have happened since that time.

Two observations, however, seem uniquely important. Claude Bernard wrote: "Art is I; Science is We." I cannot emphasize too strongly this statement in reflecting over the history of the Laboratory of Computer Science which it has been my privilege to lead for almost twenty-five years. I have been enormously fortunate in having had the opportunity to work with a large number of enormously creative, imaginative, and effective colleagues who have played a major role in the contributions we have made in the past twenty years. No presentation would be complete and accurate without giving recognition to these individuals. I have often felt that my primary responsibility was to recruit good colleagues, and then try to provide the environment where they could be optimally productive.

There are too many people to name individually in this short presentation, but I will identify some of the key individuals in relation to specific projects. I know that I will fail to recognize many individuals who have made important contributions; this failure haunts me in my writing of an historical account of the Laboratory.

The second observation I would offer is that much of what is accomplished in life and in science seems in large part to be dominated by luck, by the accident of being in the right place at the right time, with the right resources, the right funding, the right opportunities. I believe this to be true in much of what we have accomplished at the Laboratory. I feel most fortunate that thirty years ago, due to a chance association of living together with a group of MIT electrical engineering graduate students, I became acquainted with computer technology and became captivated by it as an idle recreation. I still live in fear that someday it will be discovered that I am getting paid for doing my hobby, and that I will have to get a real job.

Preparing this manuscript has been somewhat of a revelation and somewhat distressing. Technology has developed at a fantastic rate and yet progress in implementation has been laborious and slow. I found it interesting and somewhat entertaining to read two paragraphs I wrote in reviewing this field for the New England of Medicine [1] almost twenty years ago:

"Early interest in bringing the revolution in computer technology to bear on medical practice was plagued with over enthusiasm, naiveté and unrealistic expectations. The use of computers would, it was held, allow rapid and accurate collection and retrieval of all clinical information, perform automatic diagnosis, collect, monitor and analyze a variety of physiological signals, perform and interpret all laboratory tests immediately, and replace the telephone and the medical record by fulfilling their functions. In fact, however, attempts to apply computer technology to medicine have had only limited success, with numerous failures. The growing pains encountered in applying computer technology are not unique to medicine; the same types of experiences have been realized in all other areas of computer application. Indeed, one of the benefits of the computer revolution is that it stimulates and requires an intensity of thinking, a level of sophistication and a strictness of semantic behavior that are needed badly in the development of improved methods of delivering health care."

"The initial wave of optimism and enthusiasm generated by beguiling promises of an immediately available total hospital computer system has passed. Now, efforts are directed toward the painful, slow evolutionary

process of developing and implementing modules or building blocks for individual functions. There is now a keen appreciation of the wide gap separating a demonstration project, however impressive, and an operational service system in daily use. Stringent reliability requirements and the difficulties attendant when non-technical personnel (with a high turn-over rate) use a computer system on a round-the-clock basis have been two of the key limiting factors in HIS development. In addition, the very high initial costs often stagger hospital administrators; demonstration of cost-benefit savings is difficult. If the experience of other industries is repeated in hospitals, the use of computers in hospitals will not reduce total medical-care costs, but will lead to more effective use of the resources at hand and to improved patient care."

We are still in this painful, slow evolutionary process. We still face the reliability problems, the high initial costs. In reading these old articles I am impressed at how valid were our concerns twenty years ago, and how many of the issues we confronted then are still valid. For example, let me share several paragraphs of an article published in 1969 [2]:

"In papers devoted to the problems of hospital information systems, there is often a strong feeling of deja vu. For the past decade it has been repeatedly claimed that computers will be of enormous usefulness in patient care and in hospital practice. On innumerable occasions our old men have dreamed dreams and our young men have seen visions. Yet, when we critically examine what is actually being implemented in our hospitals, we are most impressed by the number of slow or halting starts, and the number of projects that have been abandoned or in which the objectives have been greatly watered down. The discrepancies between the visions and the realities are startling.

What is the cause of this state of underachievement when the need is so great, and the technology supposedly so powerful? This is a significant problem which merits considerable analysis. It is our thesis that a considerable portion of the difficulty involves the recognition and solution of the "interface" problems.

There are three important classes of interface problems: (1) the interface of the hospital with the development and implementation of a computer system; (2) the interface of the user with the computer terminals; and (3) the interface of the programmer with the computer system.

When one reviews the developmental projects which have been undertaken, there are several recurring themes. First, the magnitude of the problem is usually grossly underestimated and there is almost always inadequate concern for defining objectives and for planning. Hospital and medical staffs have had little prior experience in innovation in the area of information processing, and in many situations, the critical decisions are made by individuals in isolated departments who do not possess a broad view of the needs of the hospital.

Second, the computer industry has often displayed a considerable lack of understanding and of sophistication. On a number of occasions, the sales and promotional aspects of marketing hospital computer systems have been quite misleading and have led to false expectations.

Third, the hospitals have rarely made the depth of commitment of both administrative and professional staff that is required to develop and implement a viable system. This commitment must be in terms of years, in order to provide the exposure and experience necessary to copy with the complexity of the problems.

Fortunately, the illusionary concept of a total hospital information system is now disappearing, and there is a wider acceptance of the need for an evolutionary step-wise development over a period of years. Total systems suffer from limitations in terms of both definition and appropriate technology. We neither understand all the dimensions of the problem, nor do we have management structures or computer facilities that could support a total system. It now appears that the most logical course is the development of a series of modular systems, which will be interconnected and will evolve toward handling more and more of the information-processing aspects of patient care. It has been repeatedly demonstrated that the introduction of a radically new technology causes fundamental changes in the habit patterns and in the desires of the hospital staff. The objectives and procedures of the computer development must also evolve as we gain experience with computer systems in daily hospital operation."

My first flashback relates to my initial contact with the Hospital Computer Project contract held by MGH and a firm in Cambridge named Bolt Beranek and Newman (BBN). This contract, funded jointly by the National Institutes of Health and the American Hospital Association, began in 1962. The effort was directly related to the genius and imagination of one of the senior managers at BBN: Jordan Baruch. BBN was active in the early development of time-sharing computer systems, and had developed one of the first time-sharing systems on a PDP-1. Jordan had a vision of the useful impact that such a system would have on the information processing needs of a hospital, and persuaded the NIH to fund the effort. NIH insisted that there be a medical

collaborator and MGH was drafted - with little knowledge or insight on the part of MGH as to what was involved. There is an ancient letter in the files from the MGH director to the NIH stating that MGH would be glad to participate with the understanding that there would be no cost to MGH - after 25 years and the expenditure of many millions of dollars, this request seems somewhat lacking in foresight.

The time-sharing technology developed by BBN was at the cutting edge of computer science at that time, and the use of time-sharing at MGH was one of the first demonstrations of the potential power of remote access to a real-time, on-line data base. Jordan's early papers and reports are still fascinating in the range of the vision and the validity of many of the objectives. From the initial formulation, the system was conceived as being interactive and conversational in real time. The dialog was to have error-checking, to have on-line help, and to be branching in response to the particular set of entries given by the user. One of the most radical of the specifications was the plan to include "standing-orders" in the computer so that it could automatically detect and respond to specific data entries or to specific values of the data base. Jordan referred to this as "even-detection" or as an always active "demon" which continually monitored the data base. This is the first plan of which I am aware to take advantage of the computer system to use pre-formulated rules and is a model of what are now used extensively as "criterion-rules" or "on-line quality assurance".

The second flashback memory relates to how I initially became involved in the MGH medical information system developments. BBN had demonstrated impressive progress in developing computer technology, but there was little progress in hospital applications, and little involvement of the medical or nursing staff. I believe that the NIH site visit committee strongly suggested to MGH that it should take a more active role in the planning and management of the Hospital Computer Project, and that MGH should provide professional leadership. Following this advice, John Knowles (who was then General Director of the MGH) and Bob Ebert (Chief of Medicine) consulted with members of the computer student section and with NIH staff. It is my impression that two of you here today - Homer Warner and Bruce Waxman - were instrumental in my being recruited in 1964. This recruitment was facilitated by the circumstance that there was no one else they could identify who was interested in academic medicine and who had any experience in working with computers. In several conversations with Knowles and Baruch, I became intrigued with their vision and with the potential to combine my hobby in computers with the practice of medicine. I am sure that none of us anticipated the problems we would encounter, the amount of time, effort and money that would be required, or the years of development that lay ahead.

My next flashback memory is of the first two years of the activities of the Laboratory of Computer Science from 1964-1966. Our initial efforts were focused on the development of an admission discharge census system, a laboratory reporting system, and a medications ordering system.

The computer system used in the initial time-sharing development was a PDP-1 computer with 16K of 18 bit words memory which supported 5 simultaneous users. This system was later upgraded to a PDP-1 D which had 48K of core memory and which was supposedly capable of supporting 64 simultaneous users.

There were no telephone line communication tariffs which could be used for data communication but we were able to obtain a tariff for direct current telegraph communication that had been established for communication along railroad tracks. The telephone company was generous in not requiring us to lay the railroad track between BBN in Cambridge and MGH in order to justify providing the 10 character/second DC communication. The terminals were standard Model 33 Teletypes which were placed on a few of the floors of MGH in 1964-1965.

All programs were written in assembler. Initially all programming was done only on the central console teletype. Programs creation and debugging would take weeks to months, and changes to a running program had to be strongly justified. One BBN report stated the criteria for revision to be: "what contribution will such revision make to the knowledge garnered by the project" and "is a new principle involved?"

Using this PDP system located 10 miles distant from MGH, we were able to provide what for those days were very impressive demonstrations of how a computer might be useful in processing medical information. We were, however, never able to do more than limited demonstrations, and were therefore never able to test any of the programs in actual operation. The major problems were unreliability of the computer system, slow speed of communication, very limited functional capabilities, and user unfriendliness of the Teletype as an input instrument. In addition, because of the limitations of programming in assembler language, any debugging or enhancements of the programs required weeks to months to accomplish.

Probably of greatest importance, it became obvious that there was an incompatibility in philosophy between the goals of BBN and the goals of the Laboratory. This difference can be illustrated by quoting one of the early

BBN reports which stated: "For now we shall ignore questions of priority, time and cost. Let us work purely with definitions and ideas". In my judgment, these definitions and ideas which had been formulated by BBN were highly creative and had great promise. The hospital, however, marched to a different drummer in being more pragmatically oriented. The research project was rapidly losing credibility in that it could not meet the challenge of being used by regular hospital staff carrying out the normal daily tasks of patient care. Demonstration programs, however impressive, and promises for the future, however grandiose, could not substitute for the reality that all our computer programs could only function in the hothouse atmosphere of parallel operation for several days to several weeks and operated by our own research staff. There was no possibility for evaluating whether or not a hospital could be usefully served by such a system, and certainly no hope that we could persuade the hospital to spend its own money to implement such a system.

The next of these random flashbacks relates to the origin of MUMPS. At one of the first NIH project reviews, J.C. Shaw was one of the site visitors. Shaw is one of the early great pioneers of programming who worked at Rand Corporation with Newell and Simon. He had developed an interpretive language known as JOSS which he discussed over lunch at the site visit. One of the MIT students who was working with us on the project decided to implement the language on an MIT system over Christmas holidays. BBN management became fascinated with the language and over the next two years implemented a number of versions of the language under the names of Telecomp and Stringcomp. These languages provided algebraic programming support and had only primitive string handling functions and no file handling capability.

At about this time, Jordan Baruch persuaded General Electric to enter the business of time-sharing computer support for the hospital industry. This new GE subsidiary, known as Medinet, had very ambitious goals - simultaneously to develop time-sharing hardware, a new language and a complete set of hospital information applications. The company never had the opportunity to either succeed or fail, since after about six months GE decided to terminate all of its computer activities including Medinet. Although Medinet was not given the opportunity to implement a working demonstration of the proposed language, known as Filecomp, the specifications for the language were most intriguing in that they included the initial structure for a hierarchical global file structure.

The main character of the next flashback chapter is one of the most imaginative and productive computer scientists it has been my privilege to know - Neil Pappalardo. I first knew Neil when he was a student of mine at MIT and did his senior thesis in my cardiovascular laboratory. He and Curt Marble, a fellow MIT student, joined me after graduation from MIT as Research Assistants studying cardiovascular system control. We were using mathematical models to analyze the feedback control of cardiovascular function and became heavily involved in the use of computers - first a PDP-4, then a PDP-7. In about 1965-1966, Neil and Curt tried to persuade me that we could develop a programming system that would support the development of medical information systems at MGH. For some months, I tried to discourage them from what I felt to be a radical and obviously unproductive activity - after all, what competency and experience did a hospital-based group have in comparison with the giants in the field who were not yet able to provide a viable, reasonably inexpensive, easy to program, time-sharing system. Neil, however ignored my guidance, as was his usual habit, proceeded with the development of MUMPS, and in a few months, had a system that was exciting and promising. To this day, I have not been able to trace who first thought of that name, or why we did not come up with a more dignified title.

The development and use of MUMPS quickly became a dominant theme of our Laboratory. Indeed, it is my impression that one of the important reasons that MUMPS evolved into a powerful and easy to use language was that the development was carried out in an environment where there was very close communication between the potential users and the system designers. In our small group, there were no system analysts, no application programmers, no systems programmers - we were all problem solvers trying to develop a working system. In addition, we had the fortunate luxury of relatively stable support - NIH was at first not overly supportive of our work in creating the language, but because their investment in the MGH project was so large, and because they had no good alternatives for application development, they provided the funding to allow us to proceed.

The technical details of the first MUMPS implementation on a PDP-7 are of interest in comparison with the technology now available. The PDP-7 had 8K of memory (18 bit words) with three Teletypes and a 256 K fixed head disk. The complete MUMPS interpreter with input/output routines and buffers occupied 4K with active user programs residing the one of the 2K partitions. A major breakthrough occurred when we obtained our first PDP-9 computer which utilized a 24 K memory. The initial MUMPS systems were constructed before we had developed the global file handling system and DEC tapes were used extensively for data storage. One of our first ambulatory medical record systems had an active data base stored on over 100 different DEC tapes - not exactly a computer operator's dream. The development of the global system and the expanded core

memory on the PDP-9 made it possible to undertake much more challenging applications. In the 24K version of the PDP-9, we were able to have 16 users running simultaneously - albeit slowly.

The years of 1967-1972 were exciting times in the Laboratory. We were able to implement operational information systems in a variety of areas and undertook initial developments in a host of challenging medical applications. The major organizing principle of our development was a "modular" approach. We argued that because of the difficulty in formulating a "total systems" plan, it was more productive to focus on identifying functional information processing units of medical activity and to provide well-defined and clearly bounded computer systems which could be explicitly integrated into the real and concrete medical needs of operating departments of the hospital. This modular approach made it easier to obtain agreement and cooperation at the departmental level, sharply limited the initial start-up costs, and simplified the cost/benefit analysis since the computer program was closely linked to one specific operations unit. An additional advantage which proved of great value was that this modular approach resulted in a "domino effect" where success in one department served as an attractive role model for development of modules for other departments. Lastly, a major advantage of the modular approach was that it facilitated flexibility and graceful evolution since we could make specific changes to one module without having a large, unforeseen or negative impact on other activities.

The first two major hospital systems which we implemented were the Chemistry Laboratory test reporting system and the Bacteriology test reporting system. The first part of the Chemistry Laboratory system was implemented in 1968 and has been operational since that time, 7 days/week, 24 hours/day with a down-time of less than 0.2% for the past 20 years. This system has been extended, patched, enhanced, modified, expanded, and marginally rewritten by dozens of programmers during the last two decades. This year, we are finally being forced to rewrite the system and discard the old programs. I doubt that there are many systems with this longevity; if there were a NIM archives for old code, I would nominate these Chemistry Laboratory Reporting programs.

Two of the early post-doctoral trainees at the Laboratory were Bob Greenes (now director of the Decisions Systems Laboratory of the Brigham and Women's Hospital) and Jerry Grossman (now President and Director of the New England Medical Center.) Bob Greenes played an active role in the early creation of MUMPS and, in addition, was the lead developer of a computer-based system using a dialogue between the physician and a computer-controlled display screen. His system for using this to collect progress notes in a hypertensive clinic was the subject of his Ph.D. thesis in Computer Science at Harvard - one of the first Ph.D.'s awarded in what we now call medical informatics. It is of historical interest that Ted Shortliffe worked in the Laboratory on this project while he was a undergraduate student at Harvard College. There were no commercially available light pens at that time, so we had to invent our own touch sensitive receptors - a series of metallic strips overlying the CRT screen. These metallic strips were part of a tuned circuit which had a frequency shift induced by a change in capacitance when touched. One annoying problem occurred in conditions of high humidity or when the users fingers were overly moist - in such circumstances the user would occasionally receive a mild shock. This Pavlovian reinforcement did tend to keep down the length of the notes that were generated with the system.

Jerry Grossman was one of the key figures in the Laboratory's first efforts in developing a computer-based ambulatory medical record. When the Harvard Community Health Plan was first planned by the Harvard Medical School Dean in the late 60's, we were asked to work together to develop a computer-based medical record system. When the plan opened its doors, our computer-based record system was in place - though not heavily loaded since there were only eight patients enrolled. It was in large part a tribute to Jerry's enormous energy and commitment that COSTAR survived through the growing pains of both HCHP and the changes in the information system. He guided COSTAR through a number of re-writes and extensions - for a time, he even acted as the Medical Record Room supervisor for HCHP. COSTAR was our first major effort to develop a system which could be made available to the community as a public domain supported product; probably the most amazing observation is that the system is still being actively used and supported and extended by over a hundred different sites. Since the code was written by over 40 different programmers over the past 20 years, you can imagine that this system sets a new record for spaghetti code.

During the entire history of the Laboratory's existence I have enjoyed participating as a faculty member in the Department of Electrical Engineering and Computer Science at MIT. I have learned much from the students whom I have met in the different courses with which I have been associated. One of the most important opportunities I had was being one of the supervisors on the Ph.D. thesis of Tony Gorry (now Vice President for Institutional Development of Baylor College of Medicine.) Tony and I were fascinated by Homer Warner's early work on using Bayes' rule in computer-aided medical diagnosis. We became interested in using mathematical models to assist in clinical decision making. Homer was very generous in sharing a large data set of clinical

findings in congenital heart disease. Tony demonstrated that a program could employ sequential decision-making to balance the risk of making a diagnosis against the cost of further testing and the value of the evidence which could be obtained. His work in 1967-68 was one of the first efforts to give explicit consideration to the potential value of information when deciding whether to collect further data. Tony's work was an important stimulus to my interest in sequential decision-making and to the potential role of computer programs in medical education. In fact, some of the most sophisticated computer-based medical education programs which we now write are still strongly influenced by my early collaboration with Tony.

The Laboratory is now almost twenty-five years old and much has changed since those early years. In the intervening years many wonderful individuals have worked on a variety of projects ranging from an automated medication system, to a number of laboratory and radiology information systems, to medical education programs, to different ambulatory medical record systems. I have been fortunate to be working in a hospital which has been very supportive and provided challenging opportunities and outstanding colleagues. I have also been fortunate in the support and encouragement that has been provided by the different funding agencies. Perhaps someday we will have a conference on what happened during the second and third decades of the development of medical information systems.

It is very tempting to become overly sentimental in recalling the history of the Laboratory. Flashbacks follow flashbacks follow flashbacks - memories become rose-colored with distance; I tend to remember the good and reinterpret the bad. However, preparing this paper does give me one opportunity which I greatly treasure, and that is to thank those of you who have been so important in providing stimulation, leadership, support, guidance, wisdom and, most of all, fellowship over the years. I have always deeply enjoyed the many different experiences we have had, the site visits to Peoria and Columbia, the study section and committee work, the reports we have worked on together. A number of you in the audience I count as personal friends, as individuals who have been vitally important in both my professional and personal life. For all that you have meant to me, I can only express my deepest appreciation, and for the years ahead, I look forward to more of the same. I hope that I can pass on to the younger generation some of the excitement and fun and fellowship we have shared together.

REFERENCES

1. Barnett GO. Computers and patient care. NEJH. 1968; 279:1321-1327.
2. Barnett GO, Greenes RA. Interface aspects of a hospital information system. Ann. NY Acad. Science. 1969; 161:756-768.
3. Barnett GO, Castleman PA. A timesharing computer system for patient-care activities. Comp Biomed Research. 1967; 1:41-50.
4. Greenes RA, Pappalardo AN, Marble CW, Barnett GO. Design and implementation of a clinical data management system. Comp Biomed Research. 1969; 2:469-485.
5. Greenes RA, Barnett GO, Klein SW, Robbins A, Prior RE. Recording, retrieval and review of medical data by physician/computer interaction. NEJM. 1970;282:307-315.
6. Grossman JH, Barnett GO, Koepsell RD, Nesson HR, Dorsey JL, Phillips RR. An automated medical record system for a prepaid group practice. JAMA. 1973; 224:1616-1621.
7. Gorry GA, Barnett GO. Sequential diagnosis by computer. JAMA. 1968; 205:849-854.