HOSPITAL INFORMATION SYSTEMS: TODAY, TOMORROW AND AFTER.

George T. Horne
The Hospital for Sick Children, Toronto, Ontario, CANADA.

Summary:

Hospital Information Systems evolved in the wake of changes in the health care field. In-house development of an integrated system at the Hospital for Sick Children in Toronto resulted in a flexible network on three HP 3000/68 mainframes. Conversion of the system to a generalized bus structure with distributed intelligent workstations is under development. Future criteria and technical perspectives for the next decade are being investigated.

Out of the past...

As they have in almost all other areas of modern human endeavor computers have established themselves in health care, particularly in hospitals, and are showing every sign of being there to stay.

Contrary to the optimistic believes of the early sixties it has not been an easy walk. Indeed, few other fields are so full of pittfalls, beartraps and crocodile swamps as this particular one.

It has been said that those who don't learn from history are destined to re-live it. In order to investigate the - hopefully - successfull future of hospital systems, let us first look at the difficult past and the not allways glorious present.

Medicine is a very old profession. However the last 50-80 years brought more change to health care than all the previous centuries. As recently as in the first decade of this century there was, beside surgery, very little a physician could really do to cure a patient. Doctors were brilliant diagnosticians and could predict the course of an illness, there was a rather small number of drugs and techniques to alleviate pain and remove problems. Hospitals provided the necessary restful environment. There was little technology and only a few laboratory tests, most of which did not require complex equipment. The hospital was an extension of the family physician's field of action. It was controlled by doctors and control, being based mainly on the medical needs of patients, was relatively simple.

All this changed in the third and fourth decade: new tools, pharmaceuticals, particularly antibiotics, and the ever increasing use of complex technology made actual treatment and cure of illnesses an every-day practice. The costs and special skills needed to use these tools focused care more and more on hospitals, which in turn became complex entities not dissimilar to industrial plants.

Financial/social issues, e.g. availability of care, health insurance, etc. and the ever increasing costs also started to complicate the financial management of institutions, and their management became more and more the job professional managers - administrators.

As soon as usable computers appeared on the scene, enthusiasts believed that medicine was just the perfect place for their use.

So, in the early sixties we saw a group of physicians and engineers actually specifying, in a round table discussion, a suitable machine. The result was, as we know, the LINC - 8, which in turn became the PDP-8, and the rest is (non-medical) history.

Two areas of applications emerged: a) analysis of diagnostic measurements, typically ECG's, heroic patient monitoring, etc and b) processing of hospital/patient data, aimed at - hopefully - better management of information and resources.

Some major computer vendors (particulary one, associated somehow with a blue colouring) spent considerable sums in both areas, with relatively limited success - probably to their great surprise.

There was a number of reasons for this: analysts and programmers started out with their pre-conceived ideas about medical practice; diagnostic models proved to be just that, models that did not necessarily apply; the process by which a physician arrives at a diagnosis is more synthesis than anlysis; and human physiology stubbornly refuses to comply with computer logic.

On the management and information handling front it didn't look much better: most honest attempts to cost-justify computerization failed. Many companies are still at it and still find it difficult.

In spite of initial difficulties, hospital systems have gained wide acceptance in the last 10 years, have become imperative for modern health care and represent a substantial market.

It is interesting to note that to this day no computer hardware vendor has been extremely successfull in hospital information systems. Most of the existing, functional systems come from specialized software suppliers. I do not believe that we will see any major changes in this respect.

As we look at today's scene we have to start differentiating between the various major components of hospital systems. To simplify matters for the purposes of this paper we may safely state that most modern, at least medium-sized hospitals have at least some of their business functions (accounts, payroll, etc.) computerized. Basically these functions do not differ from similar applications in other industries, hence we shall not dwell on them.

At the other end of the scale we have the rapidly growing family of specialized medical instruments using computer technology as part of their basic function: computerized axial tomography (CAT), digital x-ray, ultrasonography and many more. They are now in the realm of companies marketing the instruments and DP professionals in hospitals are involved only marginally.

Let us concentrate on the middle field, known under the catch-all name of HIS - Hospital Information Systems.

We shall use the example of one hospital, the Hospital for Sick Children in Toronto, look at our present system, plans for tomorrow and dreams for the next decade.

The HSC System.

Our hospital is a 680-bed paediatric facility, aimed at all levels of care, with emphasis on tertiary care.

The first attempt for computerization was launched by the hospital in 1968 with the usual business systems and a rather grandiose plan for a "total" HIS. Part of this materialized in the form of an on-line Patient Admission/Discharge system as well as batch systems for laboratory test result, Medical Records and a number of other minor applications. The generally successfull, if expensive project, initially supported by Big Blue, ended in a major disaster in 1974 when, based on a consultants recommendation, an attempt was made to move to a remote data centre and to support a total of 10 teaching hospitals in Toronto. Result: nothing worked and the costs were astronomical.

In 1975 a small 4-member team proposed to start again, on an in-house basis. None of the commercially available packages were found satisfactory for our environment and we decided to develop our own customized system. There were severe, "post-disaster" financial constraints.

We started out with an HP 3000 CX, with 128 K and two 15 Mb disks. We decided to use SPL (for efficiency) and IMAGE only.

We wrote our own Screen Handler/Run Executive. It still stands up and so far we could see no advantage in using V-3000, or similar.

The lack of funds forced us into a design that is now the main virtue of the system. We could not start out, as most traditional HIS systems do, with one all-encompassing database. Each sub-system was to have its own database and had to be able to talk to the others. We wrote a communication file handler. To our great relief HP later invented IPC ; our purpose in life is not the development of system software, but rather of applications.

An A/D/T system, laboratory result reporting and a Medical Records systems were implemented within a year. Database space was obviously quite restricted and later grew as several 120 Mb drives came on stream. In 1978 we upgraded to a Series III.

It became very rapidly obvious that the machine would not be able to handle the volume of transactions and simultaneous on-line applications for more than a handfull of terminals with anything resembling reasonable response time. We upgraded to a 3000/64 and two years later added a second 64. This set-up enabled us to implement most of the basic applications that form a hospital system as it is perceived today. Disc space grew to 4000 Mb.

The Central Patient Index, which is, and will be the cornerstone of all patient systems holds demographic, and some additional information on almost a million patients. A Soundex program enables search on name and partial name if an ID number is not available.

All Registration functions (Emergency, Outpatients, Surgery,etc.) with their respective databases work through the CPR. Plastic ID plates and all necessary documents are produced on-line.

Laboratory sub-systems, with their databases, support laboratory test processing which includes on-line instrumentation and reorting of test results to the Nursing Stations, cumulative reports and laboratory statistics.

Introduction of terminals in the Nursing Stations represented a major challenge. We had to have a reasonably flexible communication network, that would serve us for some time into the future. We opted for the use of the hospital's telephone PBX and in addition to it we are using a second, voice-over-data network also on the standard telephone wires. This provides us with the capability for two data channels and one voice line at any location, as well as dial-up, modem sharing, etc., capabilities at all locations.

The machines were upgraded to mod.68's. A third, mod.48 machine was added for accounting & other functions. Down-time, planned (backup) or otherwise became crucial. Hospitals just refuse to conform with System Supervisor manuals.

It became necessary to mirror-write our databases to ensure functionality and data integrity - after learning some bitter lessons from

hot-site situations and system crashes. We developed double-write capabilities, using IPC files and DS as well as our own logging system.
Later-on, when "Shadow/Silhouette" became available, we switched to this product, mainly for maintenance reasons.

As opposed to most computer vendors (and systems people) we were aware of the fact that nurses dislike typing - and most of the time have other things to do. We looked for ways to minimize typing. As we are using block mode, at least we didn't have to wory about the line-by line entry imposed by most other systems and 4GL's. Although we believe that a touch-screen would be an excellent tool, we found the resolution and handling of the HP 150 disappointing (and the cost prohibitive) and opted for a no-typing approach using inverse video light-bars operated by four function keys.

The use of separate databases for application systems in conjunction with Message Files has the great advantage that we can develop whole sub-systems, debug, test, and subsequently implement by just opening the "tube" while keeping the system operational for the users. This approach is also used for interfacing externaly purchased application packages, e.g. a Pharmacy system.

Fig.1 shows the evolution of various applications as it stood in 1985. The system is well received by the users - particularly since we did away with the need for signing on. By mid-'85 we approached an average 6-8 thousand transactions per hour on some 130 terminals and began to have serious response problems. A performance analysis showed that we are simply out of CPU cycles on our "production" machine. We had to spread the on-line load over two machines and lost valuable redundancy. An upgrade of our 48 to a 68 will give us another year of expansion - and time to think and plan.

Future development & plans.

The Hospital has launched a re-building program and by 1991 all patient Wards will be moved to the new wing. All rooms will be private. We have the rare opportunity to use the next five years for testing techniques and concepts that will take our hospital into the 21st century. Fortunately it seems that hospital organization and management is unlikely to change drastically. What we can expect is even more emphasis on management information needs, aimed at cost control and containment. However, in the areas of information handling, control and manipulation we will see drastic changes and growth as medicine is entering the computer age with a vengeance.

As we look at our present HIS, it does not substantially differ from a number of commercial packages - it is certainly more flexible, well integrated and has a higher level of data security, etc. but it still leaves us with a number of problems, unresolved questions and open philosophical as well as technological approaches. Some of these, in no particular order, are:

- How do we stem the river of paper ?
- Data entry vs. information retrieval; where, why, how?
- We do have plenty of data; how can we provide information?
- The storage problem; where do we put it all?
- The terminal problem; how to avoid typing.
- The signature problem; how to be safe and stay practical.
- The flexibility problem: ease of use/user control vis-a-vis data quality & integrity.

- Vendor independence, or Can you put your eggs in several baskets?
- Response: how to improve it and how fast is fast enough?
- Survival: as dependency on the system grows, how do we keep vital functions if a machine dies?

At first glance these questions do not differ from similar ones appearing in other fields, like industry and business. But in the hospital environment we usually find some circumstances making the answers less than obvious.

We do not pretend to have all the answers; we have some, some are very obviously motherhood issues and some are intended to present a challenge to the industry and a look into a still rather cloudy crystal ball. It has to be emphasized that in the medical field we should have preferably both feet planted on the ground; this author at least would not cherish the idea of making repairs to parts of his anatomy dependent on a system that has just gone "hot" and is being patched by the Response Centre.

First of all, there isn't the slightest hope for stemming the proverbial river of paper unless the HIS becomes a part of a comprehensive communication network in the hospital. The network has to include all aspects of communication: voice, data, graphics, image and video. It is also mandatory that future workstations in active areas of medical care be able to have access to all these media.

It should be possible for a physician to e.g. dictate diagnostic findings into a digitizing voice storage system accessible to other doctors via telephone; after a period of time important data from such findings should find their way into the relevant databases and eventually end up in the patient's medical record stored on optical disc for random recall.

Most existing hospital system have set as their main goal the entry of orders and retrieval of results,etc. at the Nursing Station. This seems to be a remnant of a programmer's understanding of hospital procedure. It should be stated that unless data is entered at source, we defeat the purpose of the system. The place for order entry and entry of patient and nursing data should be the bedside.

Conversely, opposite to general belief the bedside is usually not where medical and other decisions are made, but rather at the Nursing Station, the doctor's lounge or diagnostic areas such as X-Ray, etc.

These are the places where we need Information rather than just data. The users should be able to control extensively the form and format of such information, without having to rely on computer experts to do it for them.

Having stated the above, let us investigate how our system can evolve in the stated direction. Just throwing in more CPU power is clearly not enough. We have allready experienced system overload, mainly because we are forcing the machine to do both data management and terminal control.

We intend to make one machine into a "database engine" or archival machine, performing most of the database writing in the network and, under normal circumstances, virtually no terminal transactions. Several "terminal servers" would have copies of the databases. As access would be mainly for reading, we can take advantage of multithreading (Turbo-Image) and cacheing. Should a server machine fail, terminals can be temporarily switched to the archive machine, sacrificing speed but keeping operational; should the archival machine fail, existing data would still be available till alternative rescue is performed.

There are two pre-requisites for this: Software to perform the necessary file transfers between the machines and a fast enough communication system to make it work. Fortunately both are in existence in the form of Silhouette/Caress and HP-LAN respectively. We intend to test this arrangement in 1986/87, with a view to introduce a Spectrum machine into the network, probably in 1988. Our aim is to be able to support, by 1990, up to 500 terminals.

The other important step is to give users of information substantial independence, while improving visible response time. We see a viable solution in introducing intelligent workstations instead of terminals.

On the Nursing Station level this would be represented by a "Supermicro", e.g. HP 9000, Micro-VAX, Micro-EAGLE. The common,required characteristics are specified as: UNIX Operating System, a relational DB tool common to UNIX users rather than machines (e.g. MISTRESS) and the use of "C", Pascal, or a 4GL common to the 3000/Spectrum and the Workstation. Databases would be down-loaded to the workstations; this is transparent to the user. The users gain substantial independence and at the same time better response; the average number of patients per Nursing Station is only about 30 and hence the databases are small.

Also, the number of applications used is restricted, depending on the clinical speciality of the particular Ward.

With the increasing capabilities of Supermicros we may also expect to use them as local servers for terminals in patient rooms.

By adhering to some broad standards we may also achieve a measure of vendor independence on this semi-peripheral level.

Based on these philosophies we have stated our aims for the next five years as:

- Gradual conversion of the system into a bus-structure Information Network, encompassing voice and data communication, Workstations and graphics capabilities.
- Central archival database machine(s) and terminal servers. Workstations with read-only databases (relational) for user independence.
- Capability to support 500 - 700 terminal devices with adequate response and very high system reliability by 1990. Full redundancy of main databases. Major computer mainframe upgrade ("Spectrum") approx. 1987/88.
- Data entry at source, i.e. at patient level, for Orders, Nursing Notes, Q.A., bedside Test Instruments, etc., where appropriate.
- Information retrieval with maximum flexibility and under user control at Nursing Station/Ward level. Computerized Patient Chart. Potential access do digitally stored (Optical Disc) Medical Record. Similar capabilities at Clinics.
- Inclusion of additional Diagnostic and other services into the Network.
- Vendor-independent Operating System and language support at Workstation/Peripheral processor level.
- Common 4th Generation Languages for Micro- and Mainframe systems. Transition to "C" and Pascal as preferred languages on mainframes.
- Introduction of Optical Mass Storage, Voice Input/Control and other advanced technologies.

New Technologies:

It is the last point of our objectives that focuses our attention on what will be necessary and desireable in the future hospital environment. Let us now take our crystal ball and list some of the expected areas of interest, equipment, and, of course, associated problems that have to be resolved.

Mass Document Storage.

For every patient and every stay or encounter, the hospital amasses a very large amount of various data, most of it in the form of printed or written documents, but also graphic (ECG) material and x-rays. Warehouses are being filled with medical records, that have to be kept for many years. Records are being requested,sent out, searched and researched and frequently lost. There is now hope in the form of optical disc storage. One platter will typically store 1 Gigabyte, a "jukebox" device can handle 1-2 hundred platters. Documents can be stored via a digitizing scanning device and random retrieval in about 15 seconds seems to be feasible. A substantial effort would be needed to convert existing records. However, beside the obvious savings for storage space we can also see a very important qualitative change. We may stop shipping paper: for the first time it may be possible to retrieve parts of a patient record via remote terminals.

Problems to be resolved: Legal aspects, privacy & access security, combination of documents and computer generated data, and periodic housekeeping problems as arecord may be spread over more discs.

On a smaller scale, there is CD-ROM and WORM-type disc systems, that show great promise for e.g. Nursing Manuals, teaching systems, and infrequently updated lists like drug interactions etc.

Bedside input devices:

As we move data entry closer to its source, the patient, it is fairly obvious that the old-fashioned terminal just won't do. Nurses and doctors are not typists, we want to free them from clerical work,not add to it. The ideal device would be flat and could hang on the wall; it would have colour and graphic capabilities for chart retrieval; it it would be icon-driven and controlled by touch and/or voice. It has to be able to accept some form of signature. It should also be part of the voice communication system.

Problems: price, selective access security.

It is hardly surprising that major communication companies are involved in the development of such devices; they are also increasingly involved in Hospital Systems. We can probably expect more from them then from the notoriously introvert computer manufacturers, as their livelihood depends on communications. They will eventually make terminal protocols transparent and make our lives a lot easier.

We can expect that in the forseeable future 30-40% of routine laboratory tests will be performed on-line, at the bedside. This implies a plethora of quality control and interfacing problems. One more reason for protocol standards across the industry.

Expert Systems and AI:

The medical profession is by its nature conservative and cautious. It will take some time until any, even well functioning diagnostic tools will earn acceptance. The sheer notion of AI is anathema to many doctors. There is however room, in the next decade, for a variety of uses of AI systems in combination with diagnostic instruments, quality assurance, quality control in laboratories and other bio-technical areas. A broad application field for expert systems is in training, education and clinical modelling. We are presently investigating some attempts in modelling the treatment of post-operative patients.

We do expect successful use of Expert Systems in the area of laboratory order processing and Pharmacy, as warning systems (potentially diagnosis dependent) for drug/test interactions.

Image processing:

Digital image processing has revolutionized the diagnostic field. It is quite conceivable that we shall soon see the demise of x-ray films, although there are still numerous details to be worked out concerning the understanding of human vision, relative resolution, and also acceptance by physicians. From a storage point of view, we are faced with a major problem, waiting for appropriate technology. For example, the average daily output of x-rays at HSC, assuming a reasonable resolution, would require close to 6 Gigabytes of permanent storage.

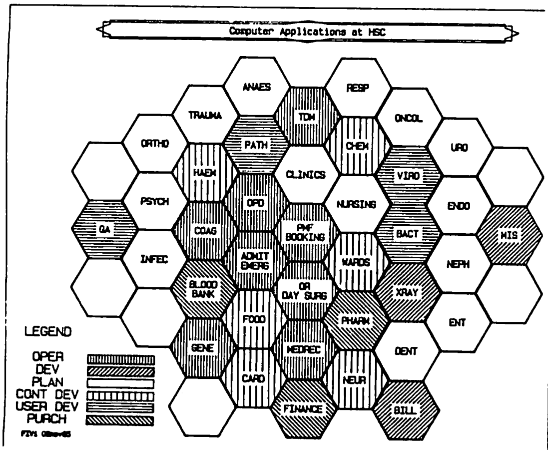

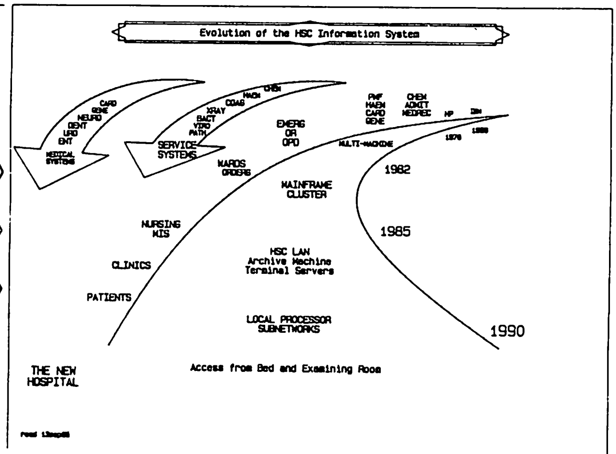

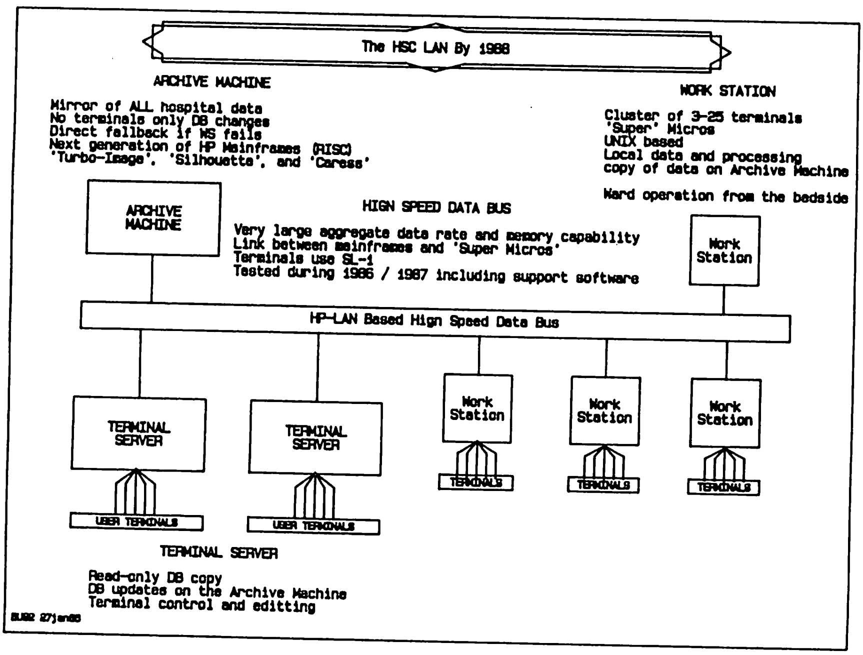

Biography

George T. Horne
has been with the Hospital for Sick Children in Toronto for 16 years, for the last 11 as Director, Computer Systems. Prior to coming to HSC he worked in medical computing at the Royal College of Surgeons of England and in continental Europe. His background and education includes Engineering (electronics) and Medicine.