The World-Wide Web past present and future, and its application to medicine.

D. M. Sendall
CERN, Geneva, Switzerland

For those interested in exploring further, a hypertext version of this paper with links to World-Wide Web sites can be found at http://www.cern.ch/CERN/WorldWideWeb/Hadrontherapy96/

Introduction

The World-Wide Web was first developed as a tool for collaboration in the high energy physics community. From there it spread rapidly to other fields, and grew to its present impressive size. As an easy way to access information, it has been a great success, and a huge number of medical applications have taken advantage of it. But there is another side to the Web, its potential as a tool for collaboration between people. Medical examples include telemedicine and teaching. New technical developments offer still greater potential in medical and other fields. This paper gives some background to the early development of the World-Wide Web, a brief overview of its present state with some examples relevant to medicine, and a look at the future.

WWW was invented and launched at CERN: how did that come about? By taking a look at the context we can see why this was a natural place to foster such a development. CERN, the European Laboratory for Particle Physics, is an international organisation with 19 member states. Its business is scientific research into the fundamental laws of matter. CERN builds and operates particle accelerators, and provides facilities for high energy physicists (particle physicists) to do their experiments. It is one of the world's largest scientific laboratories. It is also one of the oldest European joint ventures, widely regarded as an outstanding example of international co-operation.

CERN and the High Energy Physics community

Science is a community effort, and depends on free access to information and exchange of ideas. In this spirit, CERN is not an isolated laboratory, but rather a focus for a extensive community that now includes about 50 countries. Around 6000 scientists use the CERN facilities, and more than half the world's high energy physicists are now involved in its experiments. This number is increasing as CERN becomes more and more a world-wide laboratory. Although these scientists typically spend some time on the CERN site, they usually work at universities and national laboratories in their home countries.

Address for correspondence: D.M.Sendall, ECP Division, CERN, 1211 Geneva 23, Switzerland. Electronic mail Mike.Sendall@cern.ch

CERN was born out of a need to collaborate. One motive was practical, because no single European country could afford the facilities that were needed. But there was also a political dimension: to encourage co-operation between the countries of Europe that had so recently been in conflict. Pure science seemed a promising field for a joint venture, and so it has proved to be for over 40 years. The success of CERN has inspired the setting up of other international scientific organisations, such as the European Molecular Biology Laboratory EMBL, the European Southern Observatory ESO, and the European Space Agency ESA.

The way particle physics works in practice is that scientists from many countries get together to form "collaborations" with a view to setting up an experiment. Today these collaborations are large; usually they include hundreds of participating physicists, engineers and other experts. This reflects the size and complexity of the apparatus needed to carry out an experiment on the frontiers of the subject. So research teams distributed all over the world jointly design and build highly complex equipment, operate it during the running life of their experiment, and work together on the results.

Clearly, good contact must be maintained between all the people involved in such a venture. Traditional ways of staying in touch are still vital: scientists publish in learned journals, go to conferences, and exchange ideas in coffee rooms and in front of blackboards. But high energy physicists also use computers intensively to design, monitor and analyse their experiments, so it was natural to add electronic communication to the list. This is now part of everyday life, supplementing the more traditional methods.

Keeping in touch electronically

Computers and networks can be used in many ways to support the distributed activity of high energy physics. For instance, data can be sent for analysis in home institutes rather than being processed at CERN. This means that physicists can work and teach in their own institutes, keeping their home communities in touch with the latest research. Electronic distribution of software and documentation has also been standard practice for many years. The popular CERN program library is obtainable over the networks. Over the last 10-15 years, electronic mail has also become an essential tool for human communication in our field. Less formal (and faster) than a preprint or a letter, less intrusive than a telephone call, it fills a well-defined niche and improves the efficiency and ease of human communication.

CERN occupies an extensive site on the Swiss-French border near Geneva. It is served by a large and complex network of computers, with around 8000 devices interconnected by local area networks. This includes about 5000 desktop computers as well as data stores, processing centres and specialised facilities. All these are linked in turn to the rest of the world by a powerful system of networks. CERN's external links have a total capacity of around 20 million bits per second, and constitute one of the busiest Internet nodes in Europe. The laboratory is thus at the centre of an electronic web of people and institutes.

What was missing ten years ago

In spite of all this enthusiasm for electronic communication, there were many obstacles in the 1980s to the effective exchange of information. There was a great variety of computer and network systems, with hardly any common features. Users needed to understand many inconsistent and complicated systems. Different types of information had to be accessed in different ways, involving a big investment of effort by users. The result was frustration and inefficiency.

This was fertile soil for the invention of the World-Wide Web by Tim Berners-Lee. Using WWW, scientists could at last access information from any source in a consistent and simple way. The launching of this revolutionary idea was made possible by the widespread adoption of the Internet around that time [1]. This provided a de facto standard for communication between computers, on which WWW could be built. It also brought into being a "virtual community" of enthusiastic computer and communications experts, whose attitude fostered progress via the exchange of information over the Internet.

What is the World-Wide Web?

The basic idea of WWW is to merge the techniques of computer networking and hypertext into a powerful and easy to use global information system. Hypertext is text with links to further information, on the model of references in a scientific paper or cross-references is a dictionary. With electronic documents, these cross-references can be followed by a mouse-click, and with the World-Wide Web, they can be anywhere in the world.

WWW is "seamless" in the sense that a user can see the whole Web of information as one vast hypertext document. There is no need to know where information is stored, or any details of its format or organisation. Behind this apparent simplicity of course there is a set of ingenious design concepts, protocols and conventions that cannot be described here. There are many introductory and technical reference works to explain the "nuts and bolts" for those interested in learning more [2].

History and growth

The first proposal for such a system was made at CERN by Tim Berners-Lee in 1989, and further refined by him and Robert Cailliau in 1990. By the end of that year, prototype software for a basic system was already being demonstrated. To encourage the adoption of the system, it was essential to offer access to existing information without having to convert it to an unfamiliar format. This was done by providing an interface to the CERN Computer Centre's documentation and help service, and also to the familiar Usenet newsgroups. All this information immediately became accessible via a simple WWW browser, which could be run on any system.

The early system included this browser, along with an information server and a library implementing the essential functions for developers to build their own software. This was released in 1991 to the high energy physics community via the

CERN program library, so that a whole range of universities and research laboratories could start to use it. A little later it was made generally available via the Internet, especially to the community of people working on hypertext systems. By the beginning of 1993 there were around 50 known information servers.

At this stage, there were essentially only two kinds of browser. One was the original development version, very sophisticated but only available on NeXT machines. The other was the "line-mode" browser, which was easy to install and run on any platform but limited in power and user-friendliness. It was clear that the small team at CERN could not do all the work needed to develop the system further, so Tim Berners-Lee launched a plea via the Internet for other developers to join in.

Early in 1993, the National Center for Supercomputing Applications (NCSA) at the University of Illinois released a first version of their Mosaic browser [3]. This software ran in the X Window System environment, popular in the research community. It could thus offer friendly window-based interaction on a platform in widespread use. Shortly afterwards NCSA also released versions for the PC and Macintosh environments. The existence of reliable user-friendly browsers on these popular computers had an immediate impact on the spread of WWW. By late 1993 there were over 500 known servers, and WWW accounted for 1% of Internet traffic [2], which seemed a lot in those days!

Where are we now?

From these beginnings, the Web has grown into the huge system we see today. It is now so big that reliable statistics are hard to find, and they quickly become unreliable as the situation evolves [4]. At a rough estimate, it is generally thought that there are now more than 275 000 public information servers, representing one in 40 of all computers connected to the Internet. Via these servers, users can access over 30 million pages of information. Recently, many "Intranets" have also grown up, using WWW technology for communication within institutions but not making their information publicly available.

It is even harder to estimate the number of users, but recently it was suggested that there are more than 20 million in North America alone. WWW now represents a substantial fraction of Internet traffic; the last reliable statistics were for Spring 1995, when the figure was 25%. Now it may be as high as 60%, which gives rise to concern for the underlying infrastructure if expansion continues.

A success story

So it might seem that the dream has already been realised. Resources from all around the world are at your fingertips. Text, video, sound-tracks or photos were rapidly incorporated in the basic framework and are now the everyday commerce of the Web. The universe of information is inside the box that sits on your desk.

In 1945 in a now famous paper [5], Vannevar Bush foresaw:

> "a future device for individual use, which is a sort of mechanized private file and library. It ... is a device in which an individual stores all his books, records, and

communications, and which is mechanized so that it may be consulted with exceeding speed and flexibility".

He called his device the Memex: it would seem that in the form of WWW the Memex has finally been built. In fact what we have is even better, since WWW is a "virtual Memex"; the information is not actually inside the box, which can therefore be an affordable desktop computer. The information can stay where it is generated; it does not have to be laboriously transcribed and installed, and it can be kept up to date by the author or publisher. This is clearly a cheaper and more practical solution to the problem than the one Bush envisaged. But there are many ways of seeing the Web.

A reference library

The Memex model sees the Web is as a big world-wide hypertext document with easy browsing. This is just like using a reference library; the reader can browse around and follow references from a book or article along a whole trail of related material. In addition, users can add their own material; WWW has narrowed the gap between using and providing information. It is much easier and cheaper to make information available on WWW than it is to produce a reference book or journal for a library, still less to make a television programme. But it typically needs more computing skill and more investment to put information on the Web than to read what other people have provided.

Still, this modest barrier to making information available has not stood in the way of information providers. WWW as it stands has met a real need, as shown by the enormous explosion in just this "reference library" approach. Like everyone else, particle physicists use the Web heavily in this way. Today they can get access to a tutorial in Hamburg, a colleague's telephone number in Amsterdam, a preprint from California, the latest news on the top quark from Illinois, or information on how their experiment is running at CERN.

A collaborative tool for a team

This very successful approach to using the Web does not exhaust its potential. It was also intended from the first as a tool for direct human collaboration [6]. Ideally, one would like it to reproduce the environment of a few people talking around a blackboard, with access to common data. (In a medical context, these might be patient records or radiological images.) This type of application is often called "groupware" or CSCW (Computer Supported Co-operative Work). It implies a quick and easy flow of information among the participants.

Tim Berners-Lee has called this "being creative together", and "building with other people within the virtual space". Ideally, there should be no distinction between providing information and getting access to it, or between tools for browsing and for authoring. So far this goal has not been convincingly achieved with commercial software, and genuinely collaborative applications are still in their infancy. Many interesting projects are under study, but better tools are needed before they can become a reality.

What has WWW to offer medicine?

Here we can look at only a few medical examples, but they will serve to illustrate the extent of Web use today and some of its future potential. Take first the "reference library" model: there is a wealth of information available today for doctors, students, patients and many other communities. The Virtual Hospital [7] at the University of Iowa is a typical example, offering a continuously updated digital health sciences library via WWW.

Many medical journals are available on the Web, either in full versions or in the form of summaries with searchable abstracts. Clinicians and students can also consult a variety of atlases of radiology, pathology etc. The Web can offer widespread shared access to resources whose creation demands a substantial investment of time and effort. Two striking examples are the Visible Human Project [8] of the US National Library of Medicine and the consensus statements [9] of the US National Institutes of Health (NIH). Medical personnel can also easily find news of conferences in their field, or even participate in employment exchanges to hire personnel or find jobs themselves.

Many topics of medical or related interest are well covered by material on the Web, including nursing, pharmacology, health physics, and public health. Molecular biologists have been especially active in exploiting WWW technology. Scientists can access protein and DNA databases, make remote searches and use sophisticated facilities from their home laboratories.

Lost in "cyberspace"

The Web is already a victim of its own success; there is now so much information available that finding what is needed can be a real problem. But help is at hand. It can take the form of "virtual libraries" collected by real people, or of automated collections made by "robots" [10]. Robots are programs that "crawl" around the Web, visiting sites and collecting information about the material available. They use this to build indexes that can then be searched by a user query. In this way, someone with very little idea of what is available can type in a few keywords and get back a list of sites and pages with relevant information.

There are many resource guides specially tailored for medical users. They usually have links to material collected by human editors, and offer searchable indexes of their own. Health On The Internet journal reviews important medical resources and collections of resources. Examples include the Medical Imaging Internet Resources list, with links to over 100 sites, and the Health On the Net foundation which lists over 400 hospitals on the Web along with many other types of medical site. In addition to the efforts of human editors, its MARVIN medical robot accumulates a list of sites by checking for relevant medical content.

Information for patients

There are now many resources to help and inform the patient. Take the example of haematological oncology: in 1989 when the Web was first proposed, information

was not easy to find. Cancer Information Services would give an excellent response when contacted by phone or fax, but they were not easy to find or follow up, and they were not well developed outside the USA.

Today by contrast, there are around 20 information and support sites related to these diseases alone. In addition to Leukaemia there are specialised sites dealing with Hodgkin's and Non-Hodgkin's Lymphomas, Myeloma, Waldenström's Macroglobulinaemia, Myeloproliferative disorders etc. Some of these sites are maintained by patients, who try to help others by exchanging experience and giving pointers to information and support. Other sites are maintained by hospitals, foundations, or researchers. These specialised haematological sites are complemented by excellent general cancer sites. The best known of these are probably Oncolink, CancerGuide, Cansearch, TeleSCAN, and the facilities of the US National Cancer Institutes (NCI). Information on specific diseases is also typically available via sites intended for doctors.

What is really new about all this? After all, patient support information already existed before the World-Wide Web was thought of. Cancer helplines and pamphlets were available. Doctors could make literature searches using the NIH Medline service, either via networks or on compact disk at a medical library. Facilities like Physician Data Query (PDQ) offered information on clinical trials and other data for a fee. Usenet newsgroups, and support groups based on electronic mail, used the Internet to help people keep in touch with each other and exchange medical information. These media still have a role to play.

What the Web contributes is ease of use, universality, and a window on multimedia. It also offers relatively cheap, platform-independent access. This is the source of its popularity, and the growing number of users has encouraged more and better information providers via a feedback effect. As a result, everything mentioned in this section has essentially happened in the last three or four years: a truly amazing development. WWW technology is now poised to offer better technical tools than any alternative, thanks to Java and other mobile code systems (see below).

Telemedicine

In addition to the models of a reference library and a collaborating team there is another possibility of particular interest to medicine. The Web can offer widespread access to a scarce and valuable resource. This may be a human resource, such as a doctor or specialist whose time and experience are sought after, or it may be a specialised medical facility. This idea of making scarce resources available to a wider community is often called telemedicine, now a subject with its own extensive literature [11].

The development of telemedicine has traditionally been driven by special needs. Medicine in space was one, NASA being an early pioneer. Telemedicine has made it easier to provide attention for people in prison or in remote rural communities, where access to conventional medical or backup facilities may not be easy to arrange. Even emergency procedures can be carried out with help from remote experts. It has become clear that telemedicine can help to optimise the use of

doctors' time, encourage joint consultation, reduce costs and improve patient comfort. Consultations from home by the patient are another possibility, for which a lively demand has been expressed.

Teleradiology is one example of telemedicine, and many such projects are under way, often making use of WWW. For instance, the US National Library of Medicine sponsors a project at the Virtual Hospital [7]. Its aim is to explore the provision of "timely support services such as radiological consultation for local primary medical personnel" and to provide advanced image processing and visualisation tools for medical professionals. The sponsors

> *"hope that enabling medical professionals to use extremely sophisticated edge of the art diagnostic support services will minimise the need for their patients to be sent to large tertiary or academic centres just for these diagnostic support services".*

A second example among many is a project by Pennsylvania Blue Cross and Penn State Medical Center to enable community radiologists in rural facilities to send digitised radiology images to the Medical Center for consultation and evaluation. Closer to home, the University of Geneva is involved in several teleradiology projects [12].

A final example of telemedicine has been presented at this symposium. The RITA network [13] of the TERA project [14] links centres to CNAO, the Italian Centre for Hadrontherapy. It aims to provide remote consultation for patient selection, remote treatment planning, patient flow support and other facilities.

Teaching

Teaching is similar to telemedicine in many ways. Here too, the aim is to make scarce human or material resources more widely available. Students' access to computerised information can be complemented by dialogue with real specialists and teachers. The Web supplies many kinds of educational material, including medical course material. One of the best known resources is the Visible Human Project [8]. This is creating complete, anatomically detailed, three-dimensional representations of the human body. The current phase of the project is collecting transverse CT, MRI and cryosection images of representative cadavers at one millimetre intervals. The long-term goal of the project is *"to produce a system of knowledge structures that will transparently link visual knowledge forms to symbolic knowledge formats such as the names of body parts."* Many tools and projects already offer access to this data, a number of them based on WWW.

The Interactive Patient at Marshall University offers students an opportunity to go though a case study in a realistic way and make a diagnosis, which can then be commented on by specialists. The Glaxo Virtual anatomy project at Colorado State University is working on generating a 3D geometric database of the human body. A future goal of the project is to develop a virtual human anatomy laboratory for undergraduate instruction.

There are very many distance learning sites in medicine and in other fields. They can be useful not only for students, but as a means of continuing medical education for practitioners. The technique can offer an attractive mixture of access to the latest information with self-paced study for busy medical personnel. Virtual Hospital [7] for instance offers patient care support and distance learning to local physicians and other healthcare professionals. Another example is the "Virtual Medical Center", Martindale's health science guide with links to many resources.

A particularly powerful way to use the Web is to offer access to reference information so that a person can learn as much as possible, and then to let the prepared user enter into real dialogue with specialists. Such an approach can be used to offer help to General Practitioners.

Collaborative medical projects

When we consider collaborative projects in medicine a number of practical issues arise, because useful projects usually involve data on real patients. This raises security issues which are not yet properly resolved on the Internet, although progress is being made. There are also problems of standardisation, because different hospitals keep medical records in different formats. Also, in some cases there are problems of performance: high-quality medical images can take too long to transfer, so that collaboration in real time becomes impractical.

Nevertheless, the potential of collaborative projects is so great that attempts have been made to overcome these difficulties. One promising approach involves mixing private and public systems. The transfer of sensitive data can be restricted to a given site such as a hospital, or to a small number of authorised persons. It may also be possible to provide dedicated capacity over a restricted site for the rapid transfer of images. Such systems can then be linked to the wider public Internet at the cost of restricting the data available, or degrading the performance. Examples of this approach can be found in the RITA project [13] and in work at Geneva Hospital [12].

Integrated healthcare projects can be quite ambitious: one example is the ARTEMIS project [15] by CERC West Virginia. This project is designed to support a wide variety of collaborative transactions among health care providers, and WWW technology is a key ingredient. It aims to demonstrate:

> *"physicians treating patients using patient records and knowledge from distributed sources; primary care physicians consulting with remote specialists, facilitated with computer support for X-rays, ultrasound, voice annotations and other multimedia information; community care networks consisting of a collection of primary care and specialised care providers collaborating to meet a community's health care needs."*

At another level, a promising series of global healthcare applications has been identified in the context of the **G7** Information Society Pilot Projects. These typically involve linking existing distributed databases on an international level. One such project, "Towards a global public health network" will offer access to information on public health hazards and infectious diseases. Others will concentrate on sharing knowledge of best practice for the prevention, early diagnosis and treatment of

specific diseases. Another project explores the interconnection of major telemedicine centres around the world with a view to offering a 24 hour multilingual surveillance and emergency service. Work on standards, nomenclature, security, user-friendly access, translation and many other problems is needed to make effective use of the World-Wide Web in these applications.

Again not all of this is new: telemedicine has a long history using other media. For instance, in 1965 when the first intercontinental communications satellite Early Bird was launched, Dr. DeBakey performed an open-heart operation at The Methodist Hospital in the USA watched by staff at Geneva University Medical school in Switzerland, followed by an interactive question and answer session. Computer-assisted teaching also has a long and distinguished history. Once again, the vital contribution of WWW (and the Internet) is to remove the need for special equipment and software, and to add a global dimension.

So what comes next?

Where is this technology going, and what can we expect from it in the future? Such a question is notoriously difficult to answer in a rapidly evolving field. A recent informal estimate by François Fluckiger at CERN suggests that of eight important recent developments in networking technology, only one would have been predicted five years ago! So all we can usefully do is to draw attention to some obvious trends and work in progress

We can confidently expect that the underlying networking and infrastructure will get better and cheaper. This will follow further advances in technology, and also growing investment. Fibre optic cables and satellites can provide cheap high-capacity transmission. Techniques such as ADSL (Asymmetric Digital Subscriber Loop) make it possible for high-quality images to be transmitted over existing telephone circuits. Investment in ATM (Asynchronous Transfer Mode) technology should provide flexible high-performance analog (voice) and digital signals. Work continues on digital signal processing and compression techniques. Despite impressive achievements to date, these techniques have surely not yet reached their limits. The Internet architecture itself continues to evolve to cope with higher traffic and exploding numbers of users. The Web has changed the way the Internet is used, and has created a demand for consistent and high-quality service. This is being addressed by more advanced engineering in the underlying layers [16].

Deregulation of telecommunications in Europe is foreseen in principle for the end of 1998. We can hope that this will narrow the substantial gap between high-capacity transmission costs in the USA and Europe, and make available to European users a range of applications that are not economically feasible today. The commercial exploitation of the Internet should also induce businesses to invest in networking infrastructure from which all can benefit.

Investment and competition are likely to be strong in the cable and satellite television, telecommunications and computer industries as they compete to offer more popular services, especially in the entertainment field. The integration of television, networking and computing has already begun. Today from your

armchair you can press a button to browse a Web site whose address appears in a television advertisement! To some extent we can hope to get a "free ride" on the entertainment industry. There are good precedents for this: the cathode-ray tube, once a scientific curiosity, became cheap enough through mass-market television to permit today's computer video display units.

WWW-related issues

So much for the underlying technology: what can we expect from improvements in WWW itself? Here again, forecasting is dangerous, but we can outline some of the problems that preoccupy the W3 Consortium [17] and the software companies today. Work is continuing on several aspects of the WWW architecture: this will complement and take advantage of the improvements in infrastructure mentioned above. The result should be greater reliability and efficiency, along with the removal of some irritating shortcomings.

Improvements of more direct interest to users are aimed at greater flexibility and user-friendliness. The hypertext markup language is being made more powerful, and "style sheets" will give authors and readers better control over the appearance of the material on the screen. Problems of graphics and fonts are also being addressed.

Internationalisation is a pressing issue. Internet users and the computing community today are very much English-language based. This is reflected in many areas of the technology, but a serious effort is now being made to extend the benefits of the Web to other languages and writing schemes.

A number of other social issues are also under study. One of the most important for medical applications is security. Since this is also essential for financial and commercial transactions, it has been the subject of intense research. Security and confidentiality are also important for mass market applications. Encryption schemes, digital signatures and other techniques are well advanced, and we can assume that practical solutions will be available in the near future.

Another interesting development for medicine involves access to the Web by people with various disabilities. Several approaches to the use of WWW by blind or poorly sighted people in particular show considerable promise. They include Braille translation and sound synthesisers. The W3 Consortium is working on guidelines for information providers, and on browser software to make such access easier. It also maintains a list of resources available to those with various problems.

Java and multimedia

No talk on the Web today would be complete without a mention of Java [18]. With this "mobile code" approach, not only information but also the software to process it is downloaded, in a way completely transparent to the user. This revolution has just begun, and we can expect it to go far and fast. The technique has the potential to extend the power and platform-independence of many kinds of systems, but its effect on desktop software is what interests us here. In the field of data presentation and interaction it can offer powerful tools for the next generation of telemedicine. A

striking example from the Visible Human Project [8] gives an idea of what can be achieved: the Visible Human Viewer developed at Syracuse University allows any user with a simple desktop computer to extract planar views of the Visible Human dataset. Already there are applications in interactive drug design and in many other fields. Java looks especially attractive for teaching and collaborative projects.

Interesting developments can also be expected in the area of multimedia on the Web. Today by clicking on a link you can listen to a voice or an auscultation, or look at a moving picture. As the W3 Consortium puts it:

> *"WWW has been very successful in integrating new media types like images, virtual worlds or downloadable program code".*

What is missing up to now is the synchronisation of these different media. An unsophisticated example would be a slide show on the Web with synchronised voice-over. The potential for more advanced medical applications is clear.

> *"The Web offers the unique opportunity of fully integrating audio and video with many other media".*

Technical development of data models and protocols to accommodate real-time audio and video will allow the Web to be integrated with video-conferencing and other applications. Whilst there is still far to go, exciting developments on the horizon offer the possibility of creating a true distributed multimedia environment or "Compact Disc on the Web". This holds out the prospect of new collaboration and teaching tools. But again we shall need better multimedia authoring tools. We shall also need to spread awareness and experience within the community in using these new techniques.

A word of caution

This paper has painted a very positive picture of the World-Wide Web, emphasising the immense benefits it confers. But it would be unfair to conclude without mentioning at least one of its possible negative aspects. High-technology developments could put some users at a disadvantage. In particular those with poor connectivity or less advanced equipment, in the developing world for instance, might be by-passed and become second-class citizens in the new "cyberspace". In fact it is our belief that the Web can offer enormous benefits to such users too, by reducing the isolation of scientists and health practitioners, and by giving them access to up to date information from better-endowed centres. But information providers must be aware of the need to tailor their systems for such users. Appropriate techniques exist, and have already been applied in some medical applications. The community must also work hard to provide the basic connectivity and infrastructure needed by the developing world.

Acknowledgements

References

1. Krol (ed). The Whole INTERNET User's Guide and Catalog. Sebastopol, CA USA: O'Reilly Associates, 1992.

2. Berners-Lee TJ. et al. The World-Wide Web. Comm. ACM 1994; 37,8: 76-82.

3. Schatz BR. Hardin JB. NCSA Mosaic and the World Wide Web: Global Hypermedia Protocols for the Internet. Science 1994; 265: 895-901.

4. Bray T. Measuring the Web. Computer Networks and ISDN Systems 1996; 28:993-1005.

5. Bush V. As we may think. Atlantic Monthly July 1945.

6. Berners-Lee TJ. In: World Wide Web Journal Volume 1, Issue 3. Sebastopol, CA USA: O'Reilly Associates, 1996: 3-11.

7. Sims D. WWW extends Apprentice's Assistant to global medical resource. IEEE Computer Graphics 1996; 16,3: 14-15.

8. Ackerman MJ. Accessing the Visible Human Project. D-Lib Magazine, October 1995.

9. Guidelines for the planning and management of NIH Consensus Development Conferences Online. Bethesda, MD USA: National Institutes of Health, Office of the Director, Office of Medical Applications of Research; 1993 May. Updated March 1995.

10. Koster M. Robots in the Web. ConneXions 1995; 9,4.

11. Scannell KM. et al. Telemedicine: past, present, future: January 1966 through March 1995. CBM 95-4. Bethesda, MD USA: National Library of Medicine, 1995.

12. Ratib O. Medical applications of computers and networks. These proceedings

13. Ferraris M. Risso P. Squarcia A. The RITA network: a WWW application for images and clinical data transfer. These proceedings

14. Amaldi U. Oncological hadrontherapy and the TERA project. In: Greco M. (ed) 9èmes Rencontres de Physique de la Vallée d'Aoste - Results and Perspectives in Particle Physics, La Thuile, Italy. Proceedings Frascati: INFN, 1995; 605-640.

15. Jagannathan V. et al. An Overview of the CERC ARTEMIS Project. West Virginia University, Concurrent Engineering Research Center Report 1995; CERC-TR-RN-95-002.

16. Carpenter BE. Crowcroft J. Prospects for Internet Technology. Distributed Systems Engineering Journal (in press).

17. Building an Industrial Strength Web: World Wide Web Journal Volume 1, Issue 4. Sebastopol, CA USA: O'Reilly Associates, 1996.

18. Arnold K. Gosling J. The Java programming language. Reading, MA USA: Addison-Wesley, 1996.