

Internet developments and their significance for healthcare

J S Briggs and G H Early

Healthcare Computing Group, University of Portsmouth, Milton Campus, Locksway Road,
Southsea, PO4 8JF, UK

Tel: 01705 844281

Fax: 01705 844006

Email: {Jim.Briggs, Gordon.Early}@port.ac.uk

Abstract

This paper reviews some recent developments in the technology of the Internet, and shows how they may affect the way in which healthcare is provided. Starting with a brief technical history of the Internet, the paper discusses some of the technical developments that have taken place or been proposed in recent years, and speculates on the realities of their adoption within the next five years. The paper also discusses trends in public accessibility to the Internet and the development of Internet services. Finally, the impact of the technological developments on the way in which new healthcare services may be provided is discussed. Our conclusions are that the growth rate in Internet access and the improvements in performance resulting from the new technologies will make the Internet the focus of many new healthcare developments, in particular in the areas of telemedicine and in communication between patient and healthcare professionals. Increasingly, the Internet will be used to convey more "real-time" information.

1. Introduction

This paper discusses the development of the Internet, in the past and in the future, and how these developments are relevant to the healthcare field. Future developments are of course not entirely predictable so discussion of these developments is restricted to those that might occur within a five-year period. Some of the aspects outlined here have begun to be introduced in the Internet as it currently exists. Some are logical extensions of what is already available. Other suggestions are rather more speculative in nature.

2. Historical development of the Internet

The Internet was initially developed by US academic institutions and funded as a research network by US defense agencies. They were interested in developing packet switched data networks that would be resilient under battlefield conditions when some components of the network might be knocked out by enemy attack. The initial network was built in 1969 under the name ARPANET (Advanced Research Projects Agency Network) with just 4 computers linked together. In the early 1970s it was recognised that there was a need to develop a standard Internet protocol for connecting together the diverse technologies that were being experimented with. The main characteristic of this protocol was that it should be relatively simple to implement over all these diverse technologies. This protocol is now known as **Internet Protocol (IP)** and is still in use. The data packets of IP are **datagrams** that are analogous in concept to a postal letter or a telegram. In IP each packet of data is independent of all other packets just as each postal letter is independent [1]. As the Internet (or ARPANET as it then was) developed, it moved from being constructed from a single network type (now called a sub-net) to one of multiple network technologies. The need was recognised for a consistent way of connecting together these sub-nets, rather than employing ad-hoc methods of connecting each sub-net of one technology to sub-nets of differing

technologies. Quite early on in the historical development of the Internet **routers** were devised. A router performs the function of connecting different sub-nets together but, as its name suggests, its main function is to determine the path from router to router, and eventually to the destination, over which an IP datagram travels in the Internet. A given router receives an IP datagram, looks up in its routing table where to send that datagram, and when the address of the next router in the path to the destination has been found it sends the datagram to it [2]. An Internet formed out of sub-nets and connected by routers proved to be an ideal architecture to employ when the requirement arose in the mid 1980s to connect local area networks that provided site-wide communication to wide area networks that linked together different sites belonging to a company. Sometimes the Internet itself was used to link company sites but on other occasions Internet technology was employed in the creation of private corporate networks which have become known as **intranets** [3]. When the intranet is extended over the Internet via secure connections to other organisations, the resultant network is referred to as an **extranet**.

A problem with the IP datagram protocol is that it can lose data in transit over the Internet or deliver it in a different order from the order sent - both of which are unacceptable to most applications. For this reason a second standard protocol was developed in the 1970s (but not adopted for standard use until the early 1980s) called **Transmission Control Protocol (TCP)**. TCP establishes a logical connection from source to destination over which a numbered sequence of packets may flow. By identifying the position of each packet in the sequence, it is possible to put delivered packets into the correct order at the destination and ask for re-transmission of any packets that the IP protocol has lost. TCP has continued in use until today as the protocol for reliable delivery of data.

The Internet fairly rapidly became the standard means of data communication in American universities, expanding far beyond its initial goal of experimenting with resilient network

technologies. In Europe, for 10 or more years from the early 1980s, academic networks based upon the CCITT X.25 packet switched technology were employed but by the mid 1990s the Internet and Internet technology had spread world-wide.

The Internet has experienced a steady 100% growth per year over much of its life so that there are now tens of millions of computers attached to it. In the United States over 40% of households now have connections and in Europe there are perhaps 10% of households with Internet connections. Much of the recent growth has been fuelled by the increasing use of the Internet for e-mail communication and the use of the World Wide Web distributed hypertext information retrieval system for accessing information and performing transactions. The Internet is also being widely used for software distribution by file transfer and for technical support. Recent growth may not have happened had it not been for the extension of the Internet in the early 1990s from a system only accessible by companies and universities to a system that can be accessed by individual householders. This extension is not a consequence of any new technical innovation, simply the easier availability of dial-up modem connections via **Internet Service Provider (ISP)** companies.

The Internet is no longer primarily for military use or academic use but is now widely used for commercial purposes and a diverse range of uses largely based around the needs of the domestic user. It is likely that its future development will be largely fuelled by the needs of this category of user.

3. Technical developments of the Internet planned or predicted to 2003

The history of the Internet has shown its architectural design to be sufficiently flexible to permit a wide range of uses not envisaged by its original designers. We look at some of the new technical developments that build on top of that flexibility to provide more effective service. Discussion of the developments is covered under two headings:

- *Internet infrastructure* - the physical transmission media, the constituent networks from which the Internet is constructed and the methods employed for interconnecting networks and transmission media
- *Internet protocols* - low level protocols such as TCP, IP and other related low level protocols

3.1. Internet network infrastructure developments

The characteristics of the network infrastructure control, to some extent, the class of applications that can be run over the Internet. Development of this infrastructure has been important in allowing new classes of applications to be introduced. As an example of such change the 'high speed' long distance backbone of the Internet started out using 56kbps links. Those links have been replaced several times to the point where high-speed links are now at either 45Mbps or 155Mbps [4]. The faster transmission allows more data to be sent in the same time, opening the way for example for larger graphical content to complement (or replace) smaller textual content on web pages.

The Internet is composed of a variety of interconnected sub-nets that we can classify as follows:

- **Backbone networks** - providing high speed national and international links to connect together corporate local area networks (LANs) and other networks together.
- **Metropolitan Area Networks (MANs)** - providing high-speed connection between various organisations within a single metropolis or small county area.
- **Organisational LANs** with limited geographical spread, that are connected to, and are part of, the Internet. Often access into these LANs will be limited by some form of security firewall. Behind the firewall is a corporate Intranet that uses Internet technology on a private network.

- **Internet Service Provider networks** - which provide connections into the Internet from "points of presence" into which dial-in users can call.

Some of the technical developments discussed below pertain to just one of these categories of sub-net, others to several categories.

Transmission media

The Internet backbone (and other parts of the Internet) is increasingly using optical fibre for transmission because of its virtually unlimited capacity and relatively low cost. Laboratory experiments have successfully achieved transmission at data rates of well over 10,000Mbps over optical fibre and this is thought to be nowhere near the limit of what may eventually be achievable. The limiting factor is not the transmission medium itself but the production of economic transmitters and receivers to operate at high data rates and the production of computers that can use such high-speed data.

Even where the limit on transmission of single streams of data over optical fibre is reached, it is possible to send several streams simultaneously by transmitting each stream using a different light frequency, a technique known as Wave Division Multiplexing [5].

The Internet backbone currently uses speeds in the 45Mbps to 155Mbps range. In the next five years we would expect to see these rise to speeds in the range 622Mbps to 2488Mbps. A separate Internet (Internet 2) is already under construction in the USA [6] with speeds of 622Mbps which will be widely used in the five-year period, but possibly only in the USA. There is some concern in the community that, towards the end of the five-year period, only the Internet 2 might continue to be upgraded. The existing Internet would then become a lower performance byway, even though one of the objectives of the Internet 2 project is to achieve transfer of technology developed for it to the existing Internet.

We expect to see radio, satellite, and copper cable links continuing to be used in the Internet in addition to optical fibre, but not, in the main, for its backbone links due to their inferior maximum data rates.

Asynchronous Transfer Mode (ATM) sub-nets

We have already described how the Internet is composed of interconnected sub-nets. Such sub-nets may be as simple as an optical fibre link or be a complete network. It is of course possible that entirely new sub-net types will evolve in the next five years but given the increasing development costs it is unlikely that any such new sub-net types will have become widely deployed operationally. We are therefore looking at existing sub-net types for operational deployment for the backbone.

Apart from point-to-point optical fibre links the main types of wide area sub-net likely to be part of the Internet in the next five years are ATM networks and, to a lesser extent, frame relay networks. Frame relay is a limited performance transitional technology operating in the low megabit region. Many current ISP networks employ frame relay technology [7]. ATM is a scaleable technology able to operate as either LAN or WAN and currently operating typically at 155Mbps. ATM technology can operate at much higher speeds and is the technology chosen for Internet 2 where it will operate at 622Mbps. The standards defined so far for ATM run up to 2488Mbps. Though there is no ATM hardware currently operating at the higher end of this speed range it seems probable that hardware to operate at these speeds will become available within a five-year period.

ATM operates with fixed length packets (called "cells"). The use of fixed sized cells permits hardware switching from link to link over the network. This removes much of the delay between the receipt of a cell on one link and its sending out on another that is found in switches where the switching process is implemented in software.

ATM is a connection-oriented technology where a path across the ATM network is established between source and destination before any data is sent. The use of connection oriented networks provides a degree of reliability in the delivery of data and reduces the variability of delay in comparison with connectionless networks. Such technology is important where delay sensitive applications such as video and voice are to be carried by the network [8].

In many parts of the world ATM sub-nets already form a high percentage of the Internet backbone. It is expected that by the end of the five-year period ATM will be almost universally adopted for the Internet backbone.

IP switching

One of the problems with the Internet is that whatever the performance of the individual sub-nets these sub-nets have traditionally been joined together by routers which introduce significant delays into the transmission path. When a router receives a datagram it looks up in a routing table the address where it is to be sent. The datagram may then be sent out over the chosen sub-net, possibly being briefly queued if another datagram is currently being transmitted. The delay in routing between reception of a datagram at a router and its onward transmission is a major problem of the Internet.

It is now recognised that though each datagram is in principle independent of others, perhaps 80% to 90% of datagrams passing across the Internet are parts of a "flow", where a sequence of datagrams all pass from the same source to the same destination. Increasingly, intelligent router/switches are being employed which can recognise flows of data. Individual datagrams that are not part of any flow are routed conventionally, but where a flow is detected an ATM connection is set up bypassing the delay encountered in the routing process. There is already

considerable use of IP switching on the Internet and this is expected to become very widespread in the next five years, reducing the delays caused by routers [9].

The use of pure routers is likely to decline both because of the use of IP switching and because of the reduction in the use of routers between LANs, a consequence of a different form of LAN switching becoming used. Routers will still, however, continue to be used wherever one type of sub-net needs to be connected to a different type. It may be that where there are common sorts of connection between one type of sub-net and another (perhaps ATM to Ethernet) some variant of IP switching might be developed in the future to overcome that bottleneck.

Metropolitan Area Networks

In recent years there has been an increasing tendency for groups of academic and other organisations to create high-speed networks for communicating within the city or region in which they are all located. These networks are termed Metropolitan Area Networks (MANs). Organisations are connected to MANs, and MANs are connected to the Internet. This introduces an extra layer to the hierarchy of sub-nets that provides higher speed communication for local off-site data traffic than can be provided by using the Internet backbone. Frequently such MANs provide sufficiently high performance to permit the early introduction of services for metropolitan use (such as video conferencing) that cannot be supported on the wider Internet. It is highly probable that commercial use of such MANs will increase in the five-year period. An example of such commercial MAN use is the development of a MAN to link the financial institutions of the City of London.

Local Area Networks

Most corporate desktop and other computers are connected to the Internet via their own site local area network (LAN). Within a five-year period we would expect to see LANs operating

at speeds of 1,000Mbps to 5,000Mbps compared with speeds today typically in the range 10Mbps to 100Mbps. An experimental 6,400Mbps LAN was demonstrated at CERN in October 1998 [10]. 1,000Mbps LANs are just being introduced at the moment [11]. A large site will normally have an ATM LAN linking a series of 1000Mbps Ethernet or other more local LANs. Should this arrangement become common it will be possible to link the wide area ATM sub-nets to the ATM local area sub-nets using IP switching with the attendant benefits outlined earlier.

3.2. Internet protocol developments

Developments of the IP protocol

One of the more significant protocol developments in recent years has been the introduction of a new version of the IP protocol known as **IPv6** [12]. IPv6 was agreed in the early 1990s but is not yet in widespread use. The delay in its adoption is significant and is a consequence of the need to adapt routers before it can be supported. Whilst protocols which only reside in end systems (computers) attached to the Internet can be changed fairly quickly, protocols like IP which rely on changes to the basic infrastructure can only be changed at the rate that the infrastructure components are changed.

IPv6 is significant in overcoming looming problems with the current IP protocol (IPv4), viz.:

- **Address space.** The rate of growth of the Internet has meant that the 32-bit address field in IPv4 datagrams is too small and there will at some stage be an insufficient number of free IPv4 addresses to allocate to those who need them. The 128-bit address field of IPv6 will overcome this problem for the foreseeable future.
- **Resource reservation.** With the increasing use of the Internet for delay sensitive applications such as the transmission of real time video and sound for Internet telephony, there is a need to be able to guarantee that there will be no delay to datagrams carrying

digitised video or sound. A source transmitter is able to establish a 'flow' across the Internet and reserve any necessary resources in the routers on the path from source to destination. IPv4 is only able to provide reserved resources by having a separate protocol RSVP (Resource ReSerVation Protocol) [13] operating in parallel with IP to carry out the reservation. IPv6 integrates the RSVP function into the IP protocol.

We expect to see IPv6 much more widely implemented on installed routers by the end of the forecast period, and where it is implemented, its benefits will quickly become noticeable. However, address space growth is not the major driver in this. Many IP addresses are allocated dynamically (e.g. when a customer dials up an Internet Service Provider) and this sharing is keeping the rise in demand for addresses to manageable levels.

Simple Network Management Protocol (SNMP)

A network management protocol enables a monitoring station to collect information from all the routers and other network components and, if necessary, to send messages to change the status of these components. Network management has several purposes but the principal ones are the detection and subsequent recovery from network faults and the monitoring of network performance.

The principal protocol in use at the moment is SNMPv1, which was devised as a short-term measure in 1988 when the growth of the Internet suggested that manual methods of network management were becoming insufficient. A more sophisticated network management protocol, SNMPv2, has now been developed, but is being implemented rather slowly despite the manifest defects in security and other matters in SNMPv1. There is a lack of optimism about the rate of progress in its adoption so network management security will remain a difficulty. The reason for the lack of optimism is that SNMPv2 can only become employed

once it is installed on every managed network device, and upgrading devices to permit installation of SNMPv2 would add significantly to their cost [14].

Other protocols

It is possible that in the five year forecast period we may see the development of a separate layer of protocol above the TCP protocol. We have already seen the start of this with the widespread use of SSL (Secure Sockets Layer) installed with web browsers to provide a secure mechanism for data transmission [15]. The more widespread use of data compression and other mechanisms for handling different representations of data in HTML, XML, PDF, GIF, JPEG and other formats might become formalised into separate layers if and when such representations become standard ways in which to exchange data.

4. Trends in public accessibility of the Internet

Much of the growth in the use of the Internet in the 1990s is a consequence of it becoming accessible to the general public. Many business and healthcare applications of the Internet are only conceivable once there is a sufficient mass of customers/patients who are Internet connected.

The majority of domestic users have up to now accessed the Internet via the telephone network, using modems to transform the digital data coming out of the computer into an audible format acceptable to the telephone network. By using increasingly sophisticated methods of data representation the speeds attainable by such means have increased from 1200 bits per second in the early 1980s to 33,600 bits per second now, with even higher speeds available to the user with a newly standardised 56Kbits per second modem [16].

Access to the Internet has been via the telephone network to an Information Service Provider and then from that Information Service Provider's network into the Internet. This is likely to

continue to be the most common method of access in the next five years though there may be many fewer Information Service Providers to choose from as large telecommunications companies such as AT&T and BT enter the field of providing Internet access. However, on the horizon we can see a much wider range of access possibilities being open to the public. It may only become clear towards the end of this period which of these technologies will continue to be available in the longer term.

4.1. Developments in static Internet access

Apart from modems, the following technologies are likely to see use in allowing the public improved access to the Internet.

Integrated Services Digital Network (ISDN)

ISDN is a digital end to end network available since 1990 in the UK. However, only in 1998 have installation charges and the cost of equipment to connect to ISDN both been at suitably low levels for domestic users to consider this as an option.

ISDN provides two 64,000-bit/second channels over a single wire into user premises permitting simultaneous Internet access and voice calls. Alternatively the two 64kbps channels may be bonded together to give a single 128kbps link into the Internet. Though single channel ISDN speeds are only slightly higher than the highest modem speeds, the essential difference is the time to establish a connection. Connecting over ISDN typically takes less than a second, whereas it normally takes 20 to 30 seconds for a pair of modems to establish connection with one another over the telephone network. This means that ISDN can, if desired, establish a connection for each datagram sent, removing the need to pay connection charges for the long period that there may often be between transmitted datagrams. We consider that ISDN will have passed its peak usage by the end of the five-year period with other higher speed technologies becoming established by then [17].

Asymmetric Digital Subscriber Link (ADSL)

ADSL is the most prominent of a range of *Digital Subscriber Link* technologies able to provide communication at multi-megabit per second data rates between the telephone exchange and the domestic or other subscriber over normal telephone cabling. ADSL was devised for the video-on-demand market, as an alternative to the delivery of videos via cable. The asymmetry is that there is a high-speed link from the exchange to the user, and lower speed from the user to the exchange. This is a reflection of the requirement of the video on demand market but the World Wide Web user has similar requirements so ADSL can be used equally well for this purpose. Technically, the speeds available over ADSL depend on the distance from user to exchange but when it is marketed a fixed speed service is likely to be offered covering all distances at speeds of at least 2Mbps. This is over 30 times faster than the single channel ISDN speeds and removes most of the delays resulting from slow dial-up links. ADSL is likely to be paid for by means of a quarterly charge with no usage component, which will make it attractive to Internet users [18].

Other xDSL technologies being developed include HDSL (High bit-rate DSL) and SDSL (Single Line DSL), both offering 2Mbps; and VDSL (Very high bit-rate DSL) offering up to 50Mbps in one direction [19]. It seems unlikely that there will be any xDSL services offered within five years other than ADSL.

Cable modems

Cable modems permit access to the Internet using a shared cable TV channel to communicate back to the central point of a cable TV system that is linked directly or indirectly to the Internet. Such channels can be used for communicating at speeds of up to 30Mbps though first generation cable modems operate only in the range of 2Mbps to 10Mbps. There are

several difficulties with the use of cable modems that make them less likely to be as successful as ADSL links in the five-year period.

- Clearly only subscribers who have a CATV link to their premises can use cable modems.
- Cable Modems share a channel for communication to the Internet with all other cable modem subscribers on the same cable loop.
- Many older types of Cable TV networks are "broadcast only", i.e. with one way communication to the subscriber. Such networks need very extensive and expensive modification to permit them to be used for communication in the opposite direction.
- Some cable TV networks have the property of opening up access to a user's internal network to others on the same neighbourhood cable loop.

Charges for cable modem use are likely be flat rate monthly or quarterly as for ADSL [20].

Digital Powerline Technology

An alternative to DSL and cable modem technologies is **Digital Powerline (DPL)**. Developed by Nortel and Norweb, it too permits the user to be permanently connected to the Internet. Data signals (at up to 1Mbps) are carried into the customer's premises via the existing electricity power lines, having been superimposed on the power supply at the neighbourhood transformer substation. The substation and its power lines are effectively a local area network. Because users are connected via a LAN, the system suffers from the same problem as cable TV networks, namely that the more users who access the Internet via DPL, the less bandwidth is available for their use [21].

Hybrid systems

Hybrid systems use different technologies for the two directions of data flow. For example, data from the Internet may be received via a high frequency terrestrial radio transmission, or transmission via a satellite, to an aerial at the subscriber premises [22]. A "broadcast only"

cable network could also be used for this purpose. To send information to the Internet, communication may be over the telephone network or ISDN. The satellite/radio communication is a signal broadcast to all premises with the broadcast datagrams containing an address to be matched with the address of the user's receiving equipment. Only if the addresses match is the data passed to the attached computer. The kind of hybrid system in limited use now offers data rates lower than those of ADSL and cable modems and may gradually decline in use once those higher performance systems come into use. We have, however, seen in the past that higher performance systems sometimes have tariffs set which discourage their use and other more primitive technologies are able to continue in use for a considerable period of time.

WebTV

WebTV doesn't introduce any new web access methods, but by integrating web access with television viewing, it does make it possible to use information on the World-Wide Web (WWW) to support information delivered on normal broadcast television channels. The WebTV hardware is a set-top box that is linked to the television as the display device, and has a built-in modem that is connected to a telephone socket to access to the Internet. To print out information, a printer may also be connected to the set-top box. In addition, there is an infrared remote control/keyboard for selecting the pages to download and for entering alphanumeric information, perhaps for e-mail [23].

Broadcast TV and web information can be linked by broadcasting hypertext links that can be used to select web pages. This provides a degree of interactivity between the web and the television. Broadcast TV and text can be seen together via a picture-in-picture facility.

The equipment produced for the US WebTV service cannot be used in the UK because of different TV standards. However, UK set-top box and modem manufacturers Pace have been

licensed to produce a WebTV box for the European market and trials of WebTV are currently underway in the UK.

The prospects for WebTV are uncertain but there are some indicators of possible success. Many TV companies have set up web sites associated with TV programmes that provide background information, plot details and other information to support the broadcast. There have been very large numbers of "hits" on some of these web sites. In the short time that WebTV has been in operation in the US, some 300,000 subscribers have subscribed to service operated by WebTV Networks Inc, a company now owned by Microsoft. We do, however, expect a WebTV service to become firmly established in the UK in the 5 year period as one aspect of several changes now underway in the television industry, such as digital broadcasting, wide screen TV and high definition TV.

WebTV's brightest prospects may, however, come from its ability to bring the Internet into the homes of people who don't have a PC. At present, PC ownership, while growing at a fantastic rate, is still a minority activity, whereas virtually every home in the developed world has a television set. Opening up the Internet to new types of user (particularly those less technically adept than current PC owners) will provide opportunities for new applications that might not be commercially viable until a larger proportion of the population can be reached. Further, the Internet will become to many people simply an extension of the well-accepted concept of a TV channel, albeit something that could provide a limitless number of new channels.

Conclusions

Whichever access system is used, users are likely to see a reduction in communications charges over the next five years similar to or greater than that which has occurred over the past five years. This is due to reductions in the basic cost of the telecommunications

infrastructure and increased competition between operators. We would not be surprised to see a 50% reduction in costs over the period though this cannot be a firm prediction.

4.2. Mobile access

Mobile access to the Internet has been possible in recent years but only for those applications such as e-mail which don't require the transfer of much information. This is because the cellular mobile phone operates in a relatively noisy environment. There is a trade off between speed of transmission and noise. To transmit reasonably reliably the speed of transmission via mobile phones has been restricted to 9600 bits per second, less than a third of the speed of transmission over fixed wired links. Even with these limitations, it has been thought possible and worthwhile to produce specialised mobile phones with built in screens and keyboards for accessing the Internet [24]. More commonly, mobile workers accessing the Internet would use laptop PCs with a credit card sized modem plugged into their PCMCIA socket with the mobile phone attached to the PC.

In the next five years a new generation of digital mobile phones will appear that will not be so noise prone because of their use of digital transmission. Many of these phones will also have slots into which users may insert credit card sized smart cards containing on them a computer and memory. A single card can be used as multiple credit cards, as a digital electronic cash purse refillable with electronic credit over the Internet, and for many other purposes. Since a smart card is programmable there is no limit in principle to its use and healthcare applications are certainly possible. The use of smart cards will become very widespread within the five-year forecast period and their use in conjunction with the Internet is also likely to quite common. A survey by Dataquest [25] estimated that 3.4 billion smart cards would be in use by 2001.

Satellite mobile communications have been available for some years but their use has been restricted by the relatively high power needed for communication with geo-stationary satellites. A number of new satellites are being launched or are planned to be launched within the five-year forecast period which will go into low Earth orbit and will thus not require nearly such high power from the mobile transmitter. By having a large number of satellites in orbit it can be guaranteed that one of them will always be above the horizon from any point on Earth. Anywhere from 60 to 800 satellites per system will be used. From an Internet access point of view the key feature of these systems is that they permit the user to transmit and receive at speeds above 1Mbps [26] whereas earlier generations of satellite systems restricted users to 64kbps or less [27].

5. Internet service developments

As well as the changes in the technologies used to access the Internet, we are likely to see changes in the nature of the programs running in the users' computers. These implement the services for which the Internet is renowned.

5.1. Internet programming

The most significant development is that the WWW is now able to have intelligent programs running in the user, or client computer. Web browser software started out by being a tool that permitted users to select pages to be retrieved. When those pages were received from the remote server (in HTML format) the browser would interpret the HTML to produce displayed pages. The WWW has evolved to permit increased flexibility of display and input of information into forms, but it also now permits downloading of programs that can be run on the user computer.

WWW browsers such as Netscape Communicator and Microsoft Internet Explorer implement the Java Virtual Machine (JVM)[28], allowing the execution of programs written in Java.

Since Java is (almost!) machine independent and widely implemented there is no need to write different versions of software for each available computer and operating system. We expect to see many important applications created in Java to permit access of information over the Internet. We also expect to see applets written in other languages, particularly Ada, run on JVMs. In the future applets to interpret different formats of data might normally be downloaded as required. Applets may also be used for verification of data input and to create a richer graphical user environment.

We expect to see widespread use of genuine distributed operating over the Internet within the forecast period - systems in which objects distributed in an unlimited number of computers can co-operate on some common task. Work has been going on for some years to define agreed framework within which such systems can be constructed. Some 700 companies have been involved in defining the Common Object Request Broker Agency system (CORBA) which will permit distributed objects written in different languages to co-operate [29]. Microsoft has defined a similar but not identical system, called DCOM.

5.2. Security

Apart from performance, the other major issue many people and organisations are concerned about is the security of data in transit across the Internet and security violations of corporate computer systems and networks facilitated by their connection to the Internet.

Data in transit

Security systems need to ensure that data in transit:

- remains private (confidentiality);
- is delivered unmodified (integrity);
- that the receiver is assured that the information comes from who it purports to be from (authentication); and

- that once the message has been received neither sender nor receiver can deny it has been sent or received (non-repudiation) [30].

A variety of protocols have been devised for achieving these goals. Possibly the more important are SSL (Secure Sockets Layer) and SET (Secure Electronic Transactions). The former comes with the web browser but provides secure transmission for other application-oriented Internet protocols by acting as a separate security layer between TCP and protocols such as HTTP, FTP and SMTP.

SET, as its name suggests, is specifically oriented towards providing security for transactions and focuses on non-repudiation and authentication. It uses a complex system under which purchasers and vendors of services and products register themselves. Once registered, they are able to provide the other party to a transaction with an electronic certificate which may be presented to the other party to the transaction to confirm their identity. There has been concern about the complexity of SET and despite its promotion by major credit card companies its use is not yet fully agreed. Security of data in transit relies upon agreed standards by all parties involved. It seems clear that the technical issues on how to provide the requisite level of security for data in transit across the Internet are resolved but agreement on which standards to employ and for what has not been reached. It is highly likely that in the five-year forecast period security standards will be agreed upon and probable that confidence in their use will have become established [31].

Controlling access

The connection of corporate and healthcare networks to the Internet is a separate issue from that of protecting data in transit across it. It is not necessary to agree upon standards in order to connect an individual organisation, all that is required is for that organisation to decide upon the level of protection that is required and the techniques to achieve that protection.

A typical approach to protecting corporate systems is to exclude direct communication between computers within the corporate network and those that are external to it. All communication from outside the corporate network to inside it and often all those from inside to outside go via a computer system set up to act as a filter or "**firewall**". Within that computer whatever checks are deemed necessary may be incorporated to check that the data is from an authorised source to an authorised destination and contains appropriate content. It is expected that use of firewalls to protect sensitive systems will increase over the next few years as corporate paranoia and legal requirements intensify [32].

6. Consequences of developments for the delivery of new healthcare services

The technical developments outlined above will, at some point during the five-year forecast period, have several consequences detailed below which will be important for the delivery of new services.

Internet development	Likely healthcare implications
<p>The world-wide growth in the number of users of the Internet will continue at a similar level to that experienced today.</p> <p>Possible reductions in the rate of growth in mature markets such as North America being balanced by the development of new markets such as the Far East.</p> <p>The number of homes in the UK which are connected to the Internet will reach the same level as has been attained in the United States today (40%).</p>	<p>More people have access to the Internet as a source of healthcare information.</p> <p>More people communicate with their healthcare advisors via the Internet.</p>
<p>The Internet will be able to properly support delay-sensitive applications with an acceptable performance level. This means in particular that the Internet phone service will become very widely used in the five-year period.</p>	<p>Imaging applications at higher resolutions become possible.</p> <p>Integration of voice and video links with existing applications.</p>
<p>The delay experienced in the transmission of non-delay sensitive data over the Internet will be significantly lower than is experienced today despite the increased number of users.</p>	<p>Near "real-time" monitoring becomes possible.</p> <p>Better quality video conferencing.</p>

Internet development	Likely healthcare implications
<p>The Internet will cease to be seen as connecting users and servers and will become a network of "devices". This will be a result of having more addresses available when IPv6 is fully implemented and in terms of the aggregate volume of data it can carry when higher performance sub-nets are implemented and as delays at routers are reduced or removed.</p>	<p>Potential of patient-monitoring devices connected directly to the Internet.</p>
<p>Data will be able to be transmitted securely over the Internet using agreed protocols.</p>	<p>Reduced risk of unauthorised access to medical information.</p>
<p>The Internet will become an acceptable alternative means of networking for many organisations to the construction of their own private networks. The need to communicate with customers and suppliers will mean that all organisations of any size will have Internet connections even where they do not use the Internet for internal organisational communications.</p>	<p>More healthcare organisations with an Internet presence.</p> <p>More healthcare professionals use the Internet as a means of communication with each other and with patients.</p> <p>The development of healthcare oriented extranets.</p> <p>A patient can choose to have a telemedicine consultation with a healthcare worker anywhere in the developed world.</p>

Internet development	Likely healthcare implications
Integration of Internet and television (e.g. WebTV)	Open up Internet applications to new users, particularly those less technically adept.

7. Comparison with other technologies

The key difference between the Internet and any alternative approach to networking is the adaptability of the Internet. The Internet does not specify the use of any particular networking technology, and thus it has been able to incorporate local area and metropolitan area networks as they have become needed. On the wide area networking front, Frame Relay and ATM technologies are now both widely used as sub-nets for the Internet even though both of these are technologies of the late 1980s and early 1990s appearing long after the Internet was first defined.

Historically **X.25 packet switching** has been important for wide area communications. For many years this was the principal data networking technology, and telecommunications companies built and operated global X.25 packet switched networks. There are four important differences between X.25 packet switched networks and the Internet:

1. X.25 is a connection-oriented network in which a connection is formed between source and destination before any data is sent. The Internet, at the Internet Protocol level, is connectionless, though a service equivalent to that of a connection-oriented network is provided by the TCP protocol. Whilst the potential for data loss is not as great in X.25 as in the Internet, the practical differences are not so great with TCP hiding the loss of data from applications by re-transmitting any missing parts.
2. The X.25 standards only specify speeds of up to 64kbps. Though it is technically possible to run X.25 networks at speeds as high as 10Mbps, the public X.25 networks used to

connect between organisations all operate at speeds conforming to the slower X.25 standard.

3. X.25 packets are much more limited in size than IP datagrams, the former being restricted (usually) to 128 bytes, the latter being up to 64 kilobytes long. In practice the difference in size is not so great as much of the Internet restricts packet size to 512 bytes to minimise the extent to which datagrams have to be fragmented as they are transmitted over sub-nets with packet size limits.
4. In public X.25 networks there is a per packet charge whilst there are currently no charges for use of the Internet [33].

X.25 is still widely used though it is declining in importance.

Frame Relay networks are now becoming widely used, particularly for connecting together corporate LANs. Frame Relay networks have much in common with X.25 in that they are connection oriented with a limited packet size, but the network leaves responsibility for error recovery and flow control to the computers attached to it. Charging for use of such networks is not usually on a per packet basis but rather on a contract regarding the user data rate which the operator guarantees to transmit without loss. Many Information Service Providers have the networks that link dial-in customers to the Internet constructed from frame relay sub-nets. Telecommunications operators such as the BT/AT&T consortium, Concert, also offer frame relay data network services [34].

The third wide-area communication technology that is available is ATM networks. As with frame relay, ATM networks are also widely employed within the Internet and available outside the Internet. ATM networks were devised to form the communications element of a high speed "broadband ISDN" service. Broadband ISDN/ATM has been viewed as technology that could satisfy the needs of all communications users. It seems scaleable from

fairly low speeds to extremely high speeds. It can provide local or wide area networking. It can provide a connection oriented or a connectionless service (though it is fundamentally a connection oriented system). It can provide services for delay sensitive applications or applications which are not delay sensitive. It is clear that ATM is becoming increasingly important in North America and elsewhere, but it is not yet clear whether its future will be as part of a broadband ISDN service or as a major sub-net technology in the Internet.

8. Conclusions

The Internet is expected to undergo significant improvements in performance and security within the five-year forecast period, which should expand its potential for uses beyond those currently available. At the same time the numbers of users is expected to continue to grow at the historical rate of about 100% per year though perhaps the growth rate will begin to decline as the number of potential users approaches saturation. Robert Zakon [35] presents a table of the growth in the number of attached hosts showing that the number of host computers attached to the Internet grew from 19.5 million in July 1997 to 36.7 million in July 1998. By the end of the five-year period the numbers of people who will have access to the Internet will be sufficiently large for the use of the Internet for communications between healthcare professionals and with their consumers to become a common practice. We need to plan now if the useful healthcare applications are to be identified and implemented by the time consumers start to demand them.

Acknowledgements

The authors would like to acknowledge the support of Dr R.G. Curry and the financial support of the UK Department of Health in producing this paper.

References

1. Feit S. *TCP/IP - Architecture, Protocols & Implementation*, McGraw-Hill, 1993
2. Huitema C., *Routing in the Internet*, Prentice-Hall, 1995
3. Minoli D., *Internet & Intranet Engineering*, McGraw-Hill, 1997
4. Internet Backbone Maps, <http://www.nthelp.com/maps.htm>, 1998
5. Flanigan B., Digging for the Long Haul, *Telecommunications*, Nov. 1998, p55
6. Internet2, <http://www.internet2.edu/html/gigapops.html>, 1998
7. Smith P., *Frame Relay - Principles & Applications*, Addison-Wesley, 1993
8. Kyas O., *ATM Networks*, Intl. Thompson Publishing, 1995
9. Seaman M., Smarter & Faster IP Connections, *BYTE*, October 1997
10. CERN, Gigabit System Network: HIPPI-6400, <http://www.cern.ch/hsi/hippi/>, 1998
11. Gigabit Ethernet Alliance, Gigabit Ethernet Overview, http://www.gigabit-ethernet.org/technology/whitepapers/gige_97/, May 1997
12. Deering S., Hinden R., *Internet Protocol Version 6 (IPv6) Specification*, RFC 2460, <ftp://ftp.isi.edu/in-notes/rfc2460.txt>, 1998
13. Braden R., Zhang L., et al, *Resource Reservation Protocol (RSVP)*, RFC 2205, <ftp://ftp.isi.edu/in-notes/rfc2205.txt>, 1997
14. Stallings W., *SNMP, SNMPv2 & CMIP- The Practical Guide to Network Management Standards*, Addison-Wesley, 1993

15. Freir A., Karlton P. & Kocher P., *The SSL Protocol Version 3 (Internet Draft)*,
<http://home.netscape.com/eng/ssl3/draft302.txt>, Nov. 1996
16. V.90 ITU-T Recommendation - a digital & analogue modem pair for use over PSTN at
data signalling rates of up to 56000 bits/s downstream & up to 33600 bits/s upstream,
http://www.itu.int/itudoc/itu-t/approved/v/v90_49801.html, 1998
17. Ronayne J., *Integrated Services Digital Networks*, 2nd ed., UCL Press, 1996
18. ADSL Forum, ADSL Tutorial - Twisted Pair Access to the Information Highway,
http://www.adsl.com/adsl_tutorial.html, 1998
19. DSL and xDSL (Digital Subscriber Line and its Variations), <http://whatis.com/dsl.htm>,
1999
20. Azzam A., *High Speed Cable Modems: Including IEEE 802.14 Standards*, McGraw-Hill,
1997
21. Naik, G., *English Inventor Puts a Spark into the Net*, Wall Street Journal Europe, 30th
June 1998
22. Bates B., *Wireless Networked Communications - Concepts, Technology and
Implementation*, McGraw-Hill, 1994
23. Hill B., *WebTV for Dummies*, 2nd ed., IDG Books, 1998
24. Nokia 9000i Communicator, http://www.nokia.com/com9000/n9000i_spec.html, 1998
25. Smart Card Forum, Smart Card Factoids,
<http://www.smartercd.com/info/more/Factoid2.htm>, 1998

26. Technical Overview of the Teledesic Network,
<http://www.teledesic.com/tech/details.html>, 1998
27. Iridium Technology, <http://www.iridium.com/english/tech/index.html>, 1998
28. Meyer J. & Downing T., *Java Virtual Machine*, O'Reilly, 1997
29. Schmidt D., Distributed Object Computing with CORBA Middleware,
<http://www.cs.wustl.edu/~schmidt/corba.html>, Oct. 1998
30. Stallings W., *Network and Internetwork Security - Principles & Practice*, Prentice-Hall,
1995
31. Secure Electronic Transactions (SET) at Visa, <http://www.visa.com/cgi-bin/vee/nt/ecom/vee/set/main.html>, 1998
32. Cheswick R. & Bellovin S., *Firewalls & Internet Security - Repelling the Wily Hacker*,
Addison-Wesley, 1994
33. Thorpe N. & Ross D., *X.25 Made Easy*, Prentice-Hall, 1992
34. Concert Frame Relay Service, http://www.concert.com/ps_fr.asp, 1998
35. Zakon R., Hobbes' Internet Timeline v3.3,
<http://info.isoc.org/guest/zakon/Internet/History/HIT.html>, 1998

Field Cod

Field Cod

Field Cod

Field Cod

Field Cod

Field Cod