Chapter 6
The Internet Revolution

Key Topics

ARPANET
TCP/IP
The Internet
The World-Wide Web
Dot Com Bubble
E-Software Development

6.1 Introduction

The origins of what has developed to become the internet and world wide web goes back to work done in the early 1940s by Vanevar Bush (Fig. 6.1). Bush was an American scientist who had done work on submarine detection for the United States Navy, He later did work with others did work on an automated network analyser to solve differential equations at MIT. He became director of the office of Scientific Research and Development which was concerned with weapons research and development. This organisation employed several thousand scientists, and was responsible for supervising the development of the atomic bomb.

Bush developed a close win-win relationship between the United States military and universities. He arranged large research funding for universities to carry out research related to the needs of the United States military. This allowed the military to benefit from the early exploitation of research results, and it also led to better facilities and laboratories in universities to carry out research. Bush initially developed close links with universities such as Harvard and Berkeley. This was the foundation for future cooperation between the universities and the United States Department of Defence, and it would lead eventually to the development of ARPANET by DARPA.

Fig. 6.1 Vanevar Bush

Bush's influence on the development of the internet is due to his visionary description of an information management system that he called the "*memex*". The memex (memory extender) is described in his famous essay "As We May Think" which was published in the Atlantic Monthly in 1945 [Bus:45]. Bush envisaged the memex as a device electronically linked to a library and able to display books and films.

Bush's article essentially describes a theoretical proto-hypertext computer system, and this influenced the subsequent development of hypertext systems. The description of a memex is described in [Bus:45]:

> Consider a future device for individual use, which is a sort of mechanized private file and library. It needs a name, and to coin one at random, "memex" will do. A memex 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. It is an enlarged intimate supplement to his memory.
>
> It consists of a desk, and while it can presumably be operated from a distance, it is primarily the piece of furniture at which he works. On the top are slanting translucent screens, on which material can be projected for convenient reading. There is a keyboard, and sets of buttons and levers. Otherwise it looks like an ordinary desk.

Bush predicted that:

> Wholly new forms of encyclopedias will appear, ready made with a mesh of associative trails running through them, ready to be dropped into the memex and there amplified..

This description motivated Ted Nelson and Douglas Engelbart to independently formulate the various ideas that would became hypertext. Tim Berners-Lee would later use hypertext as part of the development of the world-wide web.

6.2 The ARPANET

There were approximately 10,000 computers in the world in the 1960s. These computers were very expensive ($100K–$200K) and had very primitive processing power. They contained only a few thousand words of magnetic memory, and programming and debugging of these computers was difficult. Further, communication between computers was virtually non-existent.

However, several computer scientists had dreams of world wide networks of computers, where every computer around the globe is interconnected to all of the other computers in the world. For example, Licklider[1] wrote memos in the early 1960s on his concept of an intergalactic network. This concept envisaged that everyone around the globe would be interconnected and able to access programs and data at any site from anywhere.

The United States Department of Defence founded the Advance Research Projects Agency (ARPA) in the late 1950s. ARPA embraced high-risk, high-return research and laid the foundation for what became ARPANET and later the internet. Licklider became the first head of the computer research program at ARPA, which was called the Information Processing Techniques Office (IPTO). He developed close links with MIT, UCLA, and BBN Technologies[2] and started work on his vision. Various groups, including National Physical Laboratory (NPL), the RAND Corporation and MIT, commenced work on packet switching networks. The concept of packet switching was invented by Donald Davies[3] at the NPL in 1965.

The early computers had different standards for representing data and this meant that the data standard of each computer would need to be known for effective communication to take place. There was a need to establish a standard for data representation, and a United States government committee developed the ANSII (American

[1] Licklider was an early pioneer of AI and he also formulated the idea of a global computer network. He wrote his influential paper "*Man–Computer Symbiosis*" in 1960, and this paper outlined the need for simple interaction between users and computers.

[2] BBN Technologies (originally Bolt Beranek and Newman) is a research and development high technology company. It is especially famous for its work in the development of packet switching for the ARPANET and the Internet. It also did defense work for DARPA. BBN played an important part in the implementation and operation of ARPANET. The "@" sign used in an email address was a BBN innovation.

[3] Packet switching is a fast message communication system between computers. Long messages are split into packets which are then sent separately so as to minimise the risk of congestion. Davies also worked on the ACE computer (one of the earliest stored program computers) that was developed at the NPL in the UK in the late 1940s.

Standard Code for Information Interchange) in 1963. This became the first universal standard for data for computers, and it allowed machines from different manufacturers to exchange data. The standard allowed 7-bit binary strings to stand for a letter in the English alphabet, an Arabic numeral or a punctuation symbols. The use of 7 bits allowed 128 distinct characters to be represented. The development of the IBM-360 mainframe standardised the use of 8-bits for a word, and 12-bit or 36-bit words became obsolete.

The first wide-area network connection was created in 1965. It involved the connection of a computer at MIT to a computer in Santa Monica via a dedicated telephone line. This result showed that telephone lines could be used for the transfer of data although they were expensive in their use of bandwidth. The need to build a network of computers became apparent to ARPA in the mid-1960s, and this led to work commencing on the ARPANET project in 1966. The plan was to implement a packet switched network based on the theoretical work done on packet switching done at NPL and MIT. The goal was to have a network speed of 56 Kbps. ARPANET was to become the world's first packet switched network.

BBN Technologies was awarded the contract to implement the network. Two nodes were planned for the network initially and the goal was to eventually have nineteen nodes. The first two nodes were based at UCLA and SRI. The network management was performed by interconnected "Interface Message Processors" (IMPs) in front of the major computers. Each site had a team to produce the software to allow its computers and the IMP to communicate. The IMPs eventually evolved to become the network routers that are used today. The team at UCLA called itself the Network Working Group, and saw its role as developing the Network Control Protocol (NCP). This was essentially a rule book that specified how the computers on the network should communicate.

The first host to host connection was made between a computer in UCLA and a computer in SRI in late 1969. Several other nodes were added to the network until it reached its target of 19 nodes in 1971. The Network Working Group developed the telnet protocol and file transfer protocol (FTP) in 1971. The telnet program allowed the user of one computer to remotely log in to the computer of another computer. The file transfer protocol allows the user of one computer to send or receive files from another computer. A public demonstration of ARPANET was made in 1972 and it was a huge success. One of the earliest demos was that of Weizenbaum's ELIZA program. This is a famous AI program that allows a user to conduct a typed conversation with an artificially intelligent machine (psychiatrist) at MIT. The viability of packet switching as a standard for network communication had been clearly demonstrated.

Ray Tomlinson of BBN Technologies developed a program that allowed electronic mail to be sent over the ARPANET. Tomlinson developed the "user@host" convention, and this was eventually to become the standard for electronic mail in the late 1980s.

By the early 1970s over thirty institutions were connected to the ARPANET. These included consulting organisations such as BBN, Xerox, and the MITRE Corporation, and government organisations such as NASA. Bob Metacalfe developed

a wire-based local area network (LAN) at Xerox that would eventually become Ethernet in the mid-1970s.

6.3 TCP/IP

ARPA became DARPA (Defence Advance Research Projects Agency) in 1973, and it commenced work on a project connecting seven computers on four islands using a radio based network, as well as a project to establish a satellite connection between a site in Norway and in the United Kingdom. This led to a need for the interconnection of the ARPANET with other networks. The key problems were to investigate ways of achieving convergence between ARPANET, radio-based networks, and the satellite networks, as these all had different interfaces, packet sizes, and transmission rates. Therefore, there was a need for a network to network connection protocol, and its development would prove to be an important step towards developing the internet.

An international network working group (INWG) was formed in 1973. The concept of the transmission control protocol (TCP) was developed at DARPA by Bob Kahn and Vint Cerf, and they presented their ideas at an INWG meeting at the University of Sussex in England in 1974 [KaC:74]. TCP allowed cross network connections, and it began to replace the original NCP protocol used in ARPANET. However, it would take some time for the existing and new networks to adopt the TCP protocol.

TCP is a set of network standards that specify the details of how computers communicate, as well as the standards for interconnecting networks and computers. It was designed to be flexible and provides a transmission standard that deals with physical differences in host computers, routers, and networks. TCP is designed to transfer data over networks which support different packet sizes and which may sometimes lose packets. It allows the inter-networking of very different networks which then act as one network.

The new protocol standards were known as the transport control protocol (TCP) and the internet protocol (IP). TCP details how information is broken into packets and re-assembled on delivery, whereas IP is focused on sending the packet across the network. These standards allow users to send electronic mail or to transfer files electronically, without needing to concern themselves with the physical differences in the networks. TCP/IP is a family or suite of protocols, and it consists of four layers (Table 6.1):

The internet protocol (IP) is a connectionless protocol that is responsible for addressing and routing packets. It is responsible for breaking up and assembling packets, with large packets broken down into smaller packets when they are travelling through a network that supports smaller packets. A connectionless protocol means that a session is not established before data is exchanged. Packet delivery with IP is not guaranteed as packets may be lost or delivered out of sequence. An acknowledgement is not sent when data is received, and the sender or receiver is not

Table 6.1 TCP layers

Layer	Description
Network Interface Layer	This layer is responsible for formatting packets, and placing them on to the underlying network. It is equivalent to the physical and data link layers on OSI Model.
Internet Layer	This layer is responsible for network addressing. It includes the internet protocol, the address resolution protocol, and so on. It is equivalent to the network layer on the OSI Model.
Transport Layer	This layer is concerned with data transport, and is implemented by TCP and the user datagram protocol (UDP). It is equivalent to the transport and session layers in the OSI Model.
Application Layer	This layer is responsible for liaising between user applications and the transport layer. The applications include the file transfer protocol (FTP), telnet, domain naming system (DNS), and simple mail transfer program (SMTP). It is equivalent to the application and presentation layers on the OSI Model.

informed when a packet is lost or delivered out of sequence. A packet is forwarded by the router only if the router knows a route to the destination. Otherwise, it is dropped. Packets are dropped if their checksum is invalid or if their time to live is zero. The acknowledgement of packets is the responsibility of the TCP protocol. The ARPANET employed the TCP/IP protocols as a standard from 1983.

The late 1970s saw the development of newsgroups that aimed to provide information about a particular subject. A newsgroup started with a name that is appropriate with respect to the content that it is providing. Newsgroups were implemented via USENET, and were an early example of client-server architecture. A user dials in to a server with a request to forward a certain newsgroup postings; the server then "serves" the request.

6.4 Birth of the Internet

The origins of the internet can be traced to the United States government support of the ARPANET project. Computers in several United States universities were linked via packet switching, and this allowed messages to be sent between the universities that were part of the network.

The use of ARPANET was limited initially to academia and to the United States military, and in the early years there was little interest from industrial companies. However, by the mid-1980s there were over 2,000 hosts on the TCP/IP enabled network, and the ARPANET was becoming more heavily used and congested. It was decided to shut down the network by the late-1980s, and the National Science

Foundation in the United States commenced work on the NSFNET. This work commenced in the mid-1980s, and the network consisted of multiple regional networks connected to a major backbone. The original links in NSFNET were 56 Kbps but these were later updated to the faster T1 (1.544 Mbps) links. The NSFNET T1 backbone initially connected 13 sites, but this increased due to a growing interest from academic and industrial sites in the United States and from around the world. The NSF began to realize from the mid-1980s onwards that the internet had significant commercial potential.

The NSFNET backbone was upgraded with T1 links in 1988 and the internet began to become more international. Sites in Canada and several European countries were connected to the internet. DARPA formed the Computer Emergency Response Team (CERT) to deal with any emergency incidents arising during the operation of the network.

Advanced Network Services (ANS) was founded in 1991. This was an independent not-for-profit company, and it installed a new network that replaced the NSFNET T1 network. The ANSNET backbone operated over T3 (45 Mbps) links, and it different from previous networks such as ARPANET and NSFNET in that it was owned and operated by a private company rather than the United States government. The NSF decided to focus on research aspects of networks rather than the operational side. The ANSNET network was a move away from a core network such as NSFET, to a distributive network architecture operated by commercial providers such as Sprint, MCI and BBN. The network was connected by major network exchange points, termed Network Access Points (NAPs). There were over 160,000 hosts connected to the internet by the late 1980s.

The discovery of the world-wide web by Tim Berners-Lee at CERN in 1990 was a revolutionary milestone in computing. It has transformed the way that businesses operate as well as transforming the use of the internet from mainly academic (with some commercial use) to an integral part of peoples' lives. The invention of the world-wide web by Berners-Lee is described in the next section.

6.5 Birth of the World-Wide Web

The world-wide web was invented by Tim Berners-Lee (Fig. 6.2) in 1990 at CERN in Geneva, Switzerland. CERN is a key European and international center for research in the nuclear field, and several thousand physicists and scientists work there. Lee first came to CERN in 1980 for a short contract programming assignment. He came from a strong scientific background as both his parents had been involved in the programming of the Mark I computer at Manchester University in the 1950s. He graduated in physics in the mid-1970s at Oxford University in England.

One of the problems that scientists at CERN faced was that of keeping track of people, computers, documents, databases, etc. This problem was more acute due to the international nature of CERN, as the centre had many visiting scientists from overseas who spent several months there. It also had a large pool of permanent staff.

Fig. 6.2 Tim Berners-Lee
Photo courtesy of Wikipedia.

Visiting scientists used to come and go, and in the late 1980s there was no efficient and effective way to share information among scientists.

It was often desirable for a visiting scientist to obtain information or data from the CERN computers. In other cases, the scientist wished to make results of their research available to CERN in an easy manner. Berners-Lee developed a program called "Enquire" to assist with information sharing, and the program also assisted in keeping track of the work of visiting scientists. He returned to CERN in the mid-1980s to work on other projects, and he devoted part of his free time to consider solutions to the information sharing problem. This was eventually to lead to his breakthrough and his invention of the world-wide web in 1990.

Inventors tend to be influenced by existing inventions, and especially inventions that are relevant to their areas of expertise. The internet was a key existing invention, and it allowed world wide communication via electronic email, the transfer of files electronically via FTP, and newsgroups that allowed users to make postings on various topics. Another key invention that was relevant to Berners-Lee was that of hypertext. This was invented by Ted Nelson in the 1960s, and it allow links to be present in text. For example, a document such as a book contains a table of contents, an index, and a bibliography. These are all links to material that is either within the book itself or external to the book. The reader of a book is able to follow the link to obtain the internal or external information.

The other key invention that was relevant to Berners-Lee was that of the mouse. This was invented by Doug Engelbart in the 1960s, and it allowed the cursor to be steered around the screen. The major leap that Berners-Lee made was essentially a marriage of the internet, hypertext and the mouse into what has become the

world-wide web. His vision and its subsequent realisation was beneficial both to CERN and the wider world. He described the vision as follows [BL:00]:

> Suppose that all information stored on computers everywhere were linked. Program computer to create a space where everything could be linked to everything.

Berners-Lee essentially created a system to give every "page" on a computer a standard address. This standard address is called the universal resource locator and is better known by its acronym URL. Each page is accessible via the hypertext transfer protocol (HTTP), and the page is formatted with the hypertext markup language (HTML). Each page is visible using a web-browser. The key features of Berners-Lee invention are (Table 6.2):

Table 6.2 Features of world-wide web

Feature	Description
URL	Universal Resource Identifier (later renamed to Universal Resource Locator (URL). This provides a unique address code for each web page. Browsers decode the URL location to access the web page. For example, www.amazon.com uniquely identifies the Amazon.com host web site in the United States.
HTML	Hyper Text Markup Language (HTML) is used for designing the layout of web pages. It allows the formatting of pages containing hypertext links. HTML is standardized and controlled by the World Wide Web Consortium (http://www.w3.org).
HTTP	The Hypertext Transport Protocol (HTTP) allows a new web page to be accessed from the current page.
Browser	A browser is a client program that allows a user to interact with the pages and information on the world-wide web. It uses the HTTP protocol to make requests of web servers throughout the internet on behalf of the browser user. Berners-Lee developed the first web browser called the World Wide Web browser.

Lee invented the well-known terms such as URL, HTML and world-wide web, and these terms are ubiquitous today. He also wrote the first browser program, and this allowed users to access web pages throughout the world. Browsers are used to connect to remote computers over the internet, and to request, retrieve and display the web pages on the local machine. The invention of the world-wide web by Berners-Lee was a revolution in the use of the internet. Users could now surf the web: i.e., hyperlink among the millions of computers in the world and obtain information easily.

The early browsers included Gopher developed at the University of Minnesota, and Mosaic developed at the University of Illinois. These were replaced in later years by Netscape and the objective of its design was to create a graphical-user-interface browser that would be easy to use, and would gain widespread acceptance in the internet community. Initially, the Netscape browser dominated the browser market, and this remained so until Microsoft developed its own browser called Internet Explorer. Microsoft's browser would eventually come to dominate the browser market, after what became known as the browser wars. The eventual dominance of Microsoft internet explorer was controversial, and it was subject to legal investigations in the United States. The development of the graphical browsers led to the commercialisation of the world-wide web.

The world-wide web creates a space in which users can access information easily in any part of the world. This is done using only a web browser and simple web addresses. The user can then click on hyperlinks on web pages to access further relevant information that may be on an entirely different continent. Berners-Lee is now the director of the World Wide Web Consortium, and this MIT based organisation sets the software standards for the Web.

6.6 Applications of the World-Wide Web

Berners-Lee used the analogy of the market economy to describe the commercial potential of the world-wide web. He realized that the world-wide web offered the potential to conduct business in cyberspace without human interaction, rather than the traditional way of buyers and sellers coming together to do business in the market place.

> Anyone can trade with anyone else except that they do not have to go to the market square to do so

The invention of the world-wide web was announced in August 1991, and the growth of the web has been phenomenal since then. The growth has often been exponential, and exponential growth rate curves became a feature of newly formed internet companies, and their business plans. The world-wide web is revolutionary in that:

- No single organization is controlling the web.
- No single computer is controlling the web.
- Millions of computers are interconnected.
- It is an enormous market place of millions (billions) of users.
- The web is not located in one physical location.
- The web is a space and not a physical thing.

The world-wide web has been applied to many areas including:

- Travel Industry (Booking flights, train tickets, and hotels)
- E-Marketing
- Ordering books and CDs over the web (such as www.amazon.com)

- Portal sites (such as Yahoo and Hotmail)
- Recruitment Services (such as www.jobserve.com)
- Internet Banking
- On-line casinos (for gambling)
- Newspapers and News Channels
- On-line shopping and Shopping Malls

The prediction for the growth of e-commerce businesses in the early days of the world-wide web was that the new web-based economy would replace the traditional bricks and mortars companies. The expectation was that most business would be conducted over the web, with traditional enterprises losing market share and eventually going out of business. Exponential growth of the e-commerce companies was predicted, and the size of the new web economy was estimated to be trillions of U.S. dollars.

New companies were formed to exploit the opportunities of the web, and existing companies developed e-business and e-commerce strategies to successfully adapt to the challenge of the brave new world. Companies providing full e-commerce solutions were concerned with the selling of products or services over the web to either businesses or consumers. These business models are referred to as Business-to-Business (B2B) or Business-to-Consumer (B2C). E-commerce web sites have the following characteristics (Table 6.3):

Table 6.3 Characteristics of e-commerce

Feature	Description
Catalogue of products	The catalogue of products details the products available for sale and their prices.
Well designed and easy to use	This is essential as otherwise the web site will not be used.
	Usability requires that it is easy to find the relevant information on the web site and that the performance of the site is good.
Shopping carts	This is analogous to shopping carts in a supermarket.
	The electronic shopping cart contains the products that the user plans to purchase.
	The user may add or remove items from the cart.
Security	Security of credit card information is a key concern for users of the web, as users need to have confidence that their credit card details may not be intercepted.
	There are technologies such as Secure Socket Layer (SSL) that provide encryption of credit card Information. Encryption protects the privacy of the information from un-authorized access.
Payments	Once the user has completed the selection of purchases there is a check-out facility to arrange for the purchase of the goods.
	Payment for the products is generally by credit or debit card once the user has completed shopping.
Order Fulfillment/Order Enquiry	Once payment has been received the products must be delivered to the customer.

6.7 Dot Com Companies

The growth of the internet and world-wide web was exponential, and the boom led to the formation of many "new economy" businesses. The key characteristic of these new businesses was that business was largely conducted over the web as distinct from the "old economy" bricks and mortar companies. These new companies included the internet portal site "yahoo.com", the on-line book store "amazon.com", and the on-line auction site "ebay.com". Yahoo.com provides news and a range of services, and it earns most of its revenue from advertisements. Amazon.com sells books as well as a collection of consumer and electronic goods. Ebay.com brings buyers and sellers together in an on-line auction space and allows users to bid for items. Boo.com was an on-line fashion company that failed dramatically in late 1999. Pets.com was an on-line pet supplies and accessory company that lasted one year in business. Priceline.com is an on-line travel firm that offers airlines and hotels a service to sell unfilled seats or rooms cheaply. ETrade.com is an on-line share dealing company. Some of these new technology companies were successful and remain in business. Others were financial disasters due to poor business models, poor management and poor implementation of the new technology.

Some of these technology companies offered a replacement or internet version of a traditional bricks and mortar company, whereas others offered a unique business offering. For example, the technology company "ebay.com" offers an auctioneering internet site to consumers worldwide. There was no available service like this available, and the service offering was quite distinct from traditional auctioneering.

The internet portal company "yahoo.com" was founded by David Filo and Jerry Yang who were students at Stanford. It was initially a hobby for the two students, who used it as a way to keep track of their personal interests and their corresponding web sites on the internet. Their list of interests grew over time and became too long and unwieldy. Therefore, they broke their interests into a set of categories and then sub-categories, and this is the core concept of yahoo.com. The term "yahoo" may stand for "Yet another Hierarchical Officious Oracle" or possibly it may have been taken from the name of the primitive creatures in Gulliver Travels by the Irish author Jonathon Swift. There was a lot of interest in the site from other students, family and friends and a growing community of users. The founders realized that the site had commercial potential, and they incorporated it as a business in 1995. The company launched its initial public offering (IPO) one year later in April 1996, and the company was valued at $850 million (or approximately 40 times its annual sales). The headquarters of the company are in California.

Yahoo is a portal site and offers free email accounts to users, a search engine, news, shopping, entertainment, health, and so on. The company earns a lot of its revenue from advertisement (especially the click through advertisements that appear on a yahoo web page). Research indicates that only about 2% of users actually click through on these advertisements, and it is therefore questionable whether these advertisements are actually effective. The company also earns revenue from providing services for businesses. These include web hosting services, web tools, larger

mailboxes, personal advertisements for dating, and so on. The company has been very successful.

Amazon is a famous internet company founded by Jeff Bezos in 1995. It was one of the earliest companies to sell products over the internet, and it commenced business as an on-line bookstore. Its product portfolio diversified over time to include the sale of CDs, DVDs, toys, computer software and video games. The name "Amazon" was chosen to represent the volume of products in the company's catalogue, and it was therefore appropriate to name if after the world's largest river. The initial focus of the company was to build up the "Amazon" brand throughout the world. Its goal was to become the world's largest bookstore, and in order to build up market share it initially sold books at a loss by giving discounts to buyers. Its strategy in building the brand through advertisements, marketing and discounts worked very effectively, and the company is now well-established and recording healthy profits.

The company headquarters is in Seattle, Washington. The IPO of the company took place in 1997 and was highly successful. It has been one of the most successful internet companies, and it has become the largest on line bookstore in the world. Many products on the Amazon web site are reviewed to give the buyer an indication as to how good the product is. Amazon has all the essential characteristics for a successful internet company. It has a sound business model with a very large product catalogue; a well-designed web site with good searching facilities; good check out facilities and good order fulfillment. The company has also developed an associates model, and this allows the many associates of Amazon to receive a commission for purchases of products on Amazon made through the associate site. The associates maintain a link to Amazon from their web site, and when a purchaser follows the link, and purchases a product on Amazon they receive a commission from the company.

The on-line auction web site eBay was founded in California by Pierre Omidyar in 1995. Omiyar was a French–American programmer who was born in Paris. He moved to the United States with his parents, and later studied computing. He worked a programmer at Claris prior to setting up eBay. The eBay site brings buyers and sellers together, and it allows buyers to bid for items. The company earns revenue by charging a commission for each transaction, and it is a profitable company. The IPO of eBay took place in 1998 and was highly successful.

Millions of items, for example, computers, furniture, cars and collectibles are listed, bought and sold on eBay every day. Anything that is legal and does not violate the company's terms of service may be bought or sold on the site. The sellers on eBay are individuals and also large international companies who are selling their products and services. A buyer makes a bid for a product or service, and competes against several other bidders. The highest bid is successful, and payment and delivery is then arranged. The revenue earned by eBay is generated in a number of ways. There are fees to list a product and fees that are applied whenever a product is sold. The company is international with a presence in over twenty countries.

There have been a number of strange offerings on eBay. For example, one man offered one of his kidneys for auction as part of the market for human organs. Other unusual cases have been towns that have been offered for sale (as a joke). Any

product lists that violate eBay's terms of service are removed from the site as soon as the company is aware of them. The company also has a fraud prevention mechanism which allows buyers and sellers to provide feedback on each other, and to rate each other following the transaction. The feedback may be positive, negative or neutral, and comments may be included as appropriate. This offers a way to help to reduce fraud with unscrupulous sellers or buyers as they will receive negative ratings and comments.

ETrade.com was launched in 1996 and it allows investors the facility to purchase shares electronically, and it also provides financial information to clients. Priceline was founded by Jay Walker, and offered a service to airlines and hotels to sell unfilled seats or rooms cheaply. Priceline was valued at $10 billion at its IPO. This was despite the fact that unlike airlines it had no assets, and was actually selling flights at a loss.

6.7.1 Dot Com Failures

Several of companies formed during the dot com era were successful and remain in business today. Others had appalling business models or poor management and failed in a spectacular fashion. This section considers some of the dot com failures and highlights the reasons for failure.

Webvan.com was an on-line grocery business that was based in California. It delivered products to a customer's home within a 30-minute period of their choosing. The company expanded to several other cities before it went bankrupt in 2001. Most of the failings at Webvan were due to management as the business model itself was fine. None of the management had experience of the supermarket or grocery business, and the company spent far too much on its infrastructure too quickly, and it rapidly ran out of funding. Its investors had advised the company to build up an infrastructure to deliver groceries as quickly as possible, rather than developing partnerships with existing supermarkets. It built warehouses, purchased a fleet of delivery vehicles and top of the range computer infrastructure before it ran out of money.

Boo.com was founded in 1998 by Ernst Malmsten, Kajsa Leander and Patrik Hedel. It was an on-line fashion retailer based in the United Kingdom. The company went bankrupt in 2000 and it succeeded in wasting over $135 million in shareholder funds in less than 3 years. The company web site was poorly designed for its target audience, and it went against many of the accepted usability conventions of the time. The site was designed in the days before broadband, and the designers of the site assumed that its users would employ 56 K modems. The web site included the latest Java and Flash technologies, and it took several minutes to load the first page of the web site for most users for the first released version of the site. Further, the navigation of the web site was inconsistent, and changed as the user moved around the site. The net effect was that users were turned off the site, and despite extensive advertising by the company, users were not inclined to use the site.

Other reasons for failure included poor management and leadership, lack of direction, lack of communication between departments, spirally costs left unchecked, hiring staff and contractors in large numbers leading to crippling pay roll costs. Further, a large number of products were returned by purchasers, and there was no postage charge applied for this service. However, the company incurred a significant cost for covering postage for returns. Eventually, the company went bankrupt in 2000. A description of the formation and collapse of the company is available in the book Boo Hoo [MaP:02]. This book is a software development horror story, and the maturity of the software development practices in place at the company may be judged by the fact that while it had up to 18 contractor companies working to develop the web-site the developers were working without any source code control mechanism in place. The latter is basic software engineering.

Pets.com was an on-line pet supply company founded in 1998 by Greg McLemore. It sold pet accessories and supplies. It had a well-known advertisement as to why one should shop at an online pet store. The answer to this question was: "Because Pets Can't Drive!". Its mascot (the Pets.com sock puppet) was also well known. It launched its IPO in February 2000 just before the dot com collapse.

Pets.com made investments in infrastructure such as warehousing and vehicles. However, in order to break even the company needed a critical mass of customers. The company's management believed that the company needed $300 million of revenue to breakeven, and that this would take a minimum of four to five years to achieve. This time period was based on growth of internet shopping and the percentage of pet owners that shopped on the Internet. By late 2000, the company realised that it would be unable to raise further capital due to the negative sentiment with technology firms following the dot com collapse. They tried to sell the company without success. The company went into liquidation at the end of November 2000 (just nine months after the IPO). However, it did return 17 cents in the dollar to its shareholders.

Kozmo.com was an online company that promised free one-hour delivery of small consumer goods. It was founded by Joseph Park and Yong Kang in New York in 1998. The point-to-point delivery was usually done by bicycle messengers, and no delivery fee was charged. However, the business model was deeply flawed, as the point to point delivery of small goods within a 1-hour period is extremely expensive. There was no way that the company could make a profit unless it charged delivery fees. The company argued that they could make savings to offset the delivery costs as they did not require retail space. The company expanded into several cities in the United States, and raised about $280 million from investors. It had planned to launch an IPO but due to the dot com bubble this was abandoned. The company shut down in 2001.

6.7.2 Business Models

A company requires a sound business model to be successful. A business model converts a technology idea or innovation into a commercial reality, and needs to

be appropriate for the company and the market that the company is operating in. A company needs a lot more than a good idea or invention to be successful. A company with an excellent business idea but with a weak business model may fail, whereas a company with an average business idea but an excellent business model may be quite successful. Many of the business models in the dot com era were deeply flawed, and the eventual collapse of many of these companies was predictable. Chesbrough and Rosenbroom [ChR:02] have identified six key components in a business model (Table 6.4):

Table 6.4 Characteristics of business models

Constituent	Description
Value Proposition	This describes how the product or service that the company provides is a solution to a customer problem. The value that the product or service provides is described from the customer's perspective.
Market Segment	This describes the customers that will be targeted, and it may include several market segments. Each segment may have different needs.
Value Chain Structure	This describes where the company fits into the value chain. The value chain [Por:98] includes activities such as inbound logistics for receiving raw materials, operations for converting raw materials to product, outbound logistics for order-fulfillment and delivery to the customer, marketing and sales activities, and customer support and maintenance services. These activities are supported by functions for procurement, technology development, HR, sales and marketing.
Revenue Generation and Margins	This describes how revenue will be generated and includes the various revenue streams (e.g., sales, investment income, support income, subscriptions, and so on). It also describes the cost structure (employee, rent, marketing, and so on) and the target profit margins.
Position in Value Network	This involves identifying competitors and other players that can assist in delivering added value to the customer.
Competitive Strategy	This describes how the company will develop a competitive advantage to be a successful player in the market. For example, the company may plan to be more competitive on cost than other players, or it may be targeting a niche market.

6.7.3 Bubble and Burst

Netscape was founded as Mosaic Communications by Marc Andreessen and Jim Clark in 1994. It was renamed as Netscape in late 1994. The initial public offering of the company in 1995 demonstrated the incredible value of the new internet companies. The company had planned to issue the share price at $14 but decided at the last minute to issue it at $28. The share price reached $75 later that day. This was

followed by what became, in effect, the dot com bubble where there were a large number of public offerings of internet stock, and where the value of these stocks reached astronomical levels. Eventually, reality returned to the stock market when it crashed in April 2000, and share values returned to more sustainable levels.

The valuations of these companies were extraordinary. The vast majority of these companies were losing substantial sums of money, and few expected to deliver profits in the short term. Traditional forms of determining the value of a company using accounts such as the balance sheet, profit and loss account, and price to earnings ratio were ignored. Instead, investment bankers argued that there was now a new paradigm in stock valuation. This brave new world argued that the potential future earnings of the stock should be considered in determining the appropriate valuation of the stock. This was used to justify the high prices of shares in technology companies, and frenzied investors rushed to buy these over-priced and over-hyped stocks. Common sense seemed to play no role in decision making, and this is a characteristic of a bubble. The dot com bubble included features such as:

- Irrational exuberance on the part of investors
- Insatiable appetite for Internet Stocks
- Incredible greed from all parties involved
- A lack of rationality by all concerned
- No relationship between Balance Sheet/Profit and Loss and the share price
- Interest in Making Money rather than in building the business first
- Following Herd Mentality
- Questionable Decisions by Federal Reserve Chairman Alan Greenspan
- Questionable Analysis by Investment Firms
- Conflict of Interest in Investment Banks between Analysis for Investors and earning revenue from IPOs
- Market had left reality behind

However, many individuals and companies made a lot of money from the boom, although many other investors including pension funds and life assurance funds made significant losses. The investment banks typically earned 5%–7% commission on each successful IPO, and it was therefore in their interest not to question the boom too closely. Those investors that bought early and disposed early obtained a good return; however, many of those who held shares for too long incurred losses.

The full extent of the boom can be seen from Figs. 6.3 and 6.4 which show the increase in value of the Dow Jones and Nasdaq from 1995 through 2002.

The extraordinary rise of the Dow Jones from a level of 3800 in 1995 to 11900 in 2000 represented a 200% increase over 5 years or approximately 26% annual growth (compound) over 5 years. The rise of the Nasdaq over this time period is even more dramatic. It rose from a level of 751 in 1995 to 5000 in 2000 representing a 566% increase over the time period. This is equivalent to a 46% compounded annual growth rate of the index.

The fall of the indices have been equally dramatic especially in the case of the Nasdaq which peaked at 5000 in March 2000 had fallen by 76%–1200% by September 2002. It had become clear that internet companies were rapidly going through

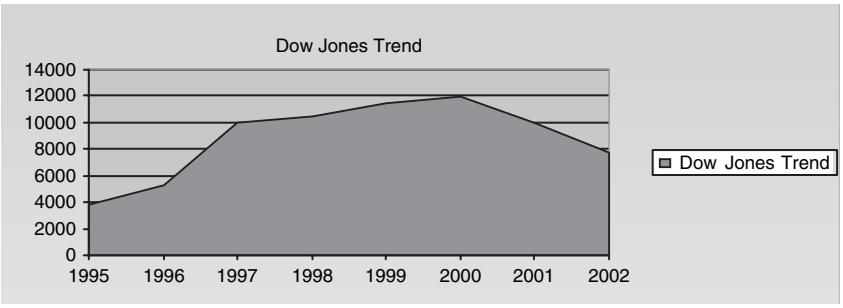

Fig. 6.3 Dow Jones (1995–2002)

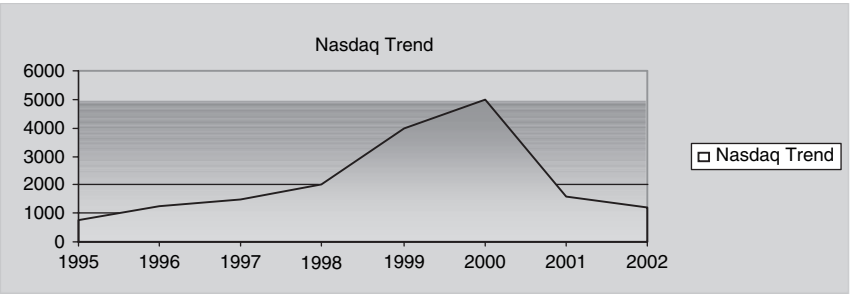

Fig. 6.4 Nasdaq (1995–2002)

the cash raised at the IPOs, and analysts noted that a significant number would be out of cash by end of 2000. Therefore, these companies would either go out of business, or would need to go back to the market for further funding. This led to questioning of the hitherto relatively unquestioned business models of many of these internet firms. Funding is easy to obtain when stock prices are rising at a rapid rate. However, when prices are static or falling, with negligible or negative business return to the investor, then funding dries up. The actions of the Federal Reserve in rising interest rates to prevent inflationary pressures also helped to correct the irrational exuberance. However, it would have been more appropriate to have taken this action 2–3 years earlier.

Some independent commentators had recognized the bubble but their comments and analysis had been largely ignored. These included "The Financial Times" and the "Economist" as well as some commentators in the investment banks. Investors rarely queried the upbeat analysis coming from Wall Street, and seemed to believe that the boom would never end. They seemed to believe that rising stock prices would be a permanent feature of the United States stock markets. Greenspan had argued that it is difficult to predict a bubble until after the event, and that even if the bubble had been identified it could not have been corrected without causing a contraction. Instead, the responsibility of the Fed according to Greenspan was to mitigate the fallout when it occurs.

There have, of course, been other stock market bubbles throughout history. For example, in the 1800s there was a rush on railway stock in England leading to a bubble and eventual burst of railway stock prices in the 1840s.

6.8 E-Software Development

The growth of the world-wide web and electronic commerce in recent years has made the quality, reliability and usability of the web sites a key concern. An organisation that conducts part or all of its business over the world-wide web will need to ensure that its web site is fit for purpose. Software development for web-based systems is relatively new, and an introduction is provided in [Nie:01]. Web applications are quite distinct from other software systems in that:

- They may be accessed from anywhere in the world.
- They may be accessed by many different browsers.
- The world-wide web is a distributed system with millions of servers and billions of users.
- The usability and look and feel of the application is a key concern.
- The performance of the web site is a key concern.
- Security threats may occur from anywhere.
- The web site must be capable of dealing with a large number of transactions at any time.
- The web site has very strict availability constraints (typically $24 \times 7 \times 365$ availability).
- The web site needs to be highly reliable.

There are various roles involved in web based software development including content providers who are responsible for providing the content on the web; web designers who are responsible for graphic design of the web site; programmers who are responsible for the implementation; and administrators who are responsible for administration of the web site. These roles may be combined in practice.

Rapid Application Development (RAD) or Joint Application Development (JAD) lifecycles are often employed for web site development. The lifecycle followed is usually spiral rather than waterfall, and it is often the case that the requirements for web-based systems will not be fully defined at project initiation. This is especially the case with usability requirements that generally evolve to the final agreed set. It is generally inappropriate to employ the waterfall lifecycle for this domain.

The software development proceeds in a number of spirals where each spiral typically involves updates to the requirements, design, code, testing, and a user review of the particular iteration or spiral.

The spiral is, in effect, a re-usable prototype and the customer examines the current iteration and provides feedback to the development team to be included in the next spiral. The approach is to partially implement the system. This leads to a better understanding of the requirements of the system and it then feeds into the next

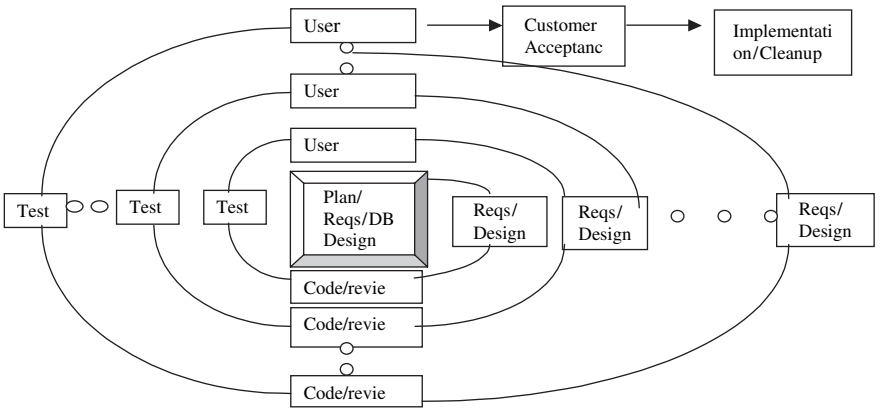

Fig. 6.5 Spiral lifecycle model

cycle in the spiral. The process repeats until the requirements and product are fully complete. The spiral model is shown in Fig. 6.5.

Web site quality and performance is a key concern to companies and their customers. The implementation of a web site requires sound software engineering practices to design, develop and test the web site.

A spiral lifecycle is usually employed and the project team needs to produce similar documentation as the waterfall model, except that the chronological sequence of delivery of the documentation is more flexible. The joint application development is important as it allows early user feedback to be received on the look and feel and correctness of the application. The approach is often "design a little, implement a little, and test a little" as this is generally the best approach to developing a usable application.

Various technologies are employed in web developed. These include HTML which is used to design simple web pages and was originally developed by Berners-Lee; CGIs (Common Gateway Interface) are often employed in sending a completed form to a server; cookies are employed to enable the server to store client specific information on client's machine. Other popular technologies are Java, Javascript, VB Script, Active X and Flash. There are tools such as Dreamweaver to assist with Web site design.

Testing plays an important role in assuring quality and various types of web testing include:

- Static testing
- Unit testing
- Functional testing
- Browser compatibility testing
- Usability testing
- Security testing
- Load/performance/stress testing

- Availability testing
- Post deployment testing

Static testing generally involves inspections and reviews of documentation. The purpose of static testing of web sites is to check the content of the web pages for accuracy, consistency, correctness, and usability, and also to identify any syntax errors or anomalies in the HTML. The purpose of unit testing is to verify that the content of the web pages correspond to the design, that the content is correct, that all the links are valid, and that the web navigation operates correctly. The purpose of functional testing is to verify that the functional requirements are satisfied. Functional testing may be quite extensive as e-commerce applications can be quite complex. It may involve product catalogue searches, order processing, credit checking and payment processing, as well as interfaces to legacy systems. Also, testing of cookies, whether enabled or disabled, needs to be considered.

The purpose of browser compatibility testing is to verify that the web browsers that are to be supported are actually supported. Different browsers implement HTML differently; for example, there are differences between the implementation by Netscape and Microsoft. These days internet explorer is the main web browser as it now dominates the market.

The purpose of usability testing is to verify that the look and feel of the application is good. The purpose of security testing is to ensure that the web site is secure. The purpose of load, performance and stress testing is to ensure that the performance of the system is within the defined parameters. There are tools to measure web server and operating system parameters and to maintain statistics. These tools allow simulation of a large number of users at one time or simulation of the sustained use of a web site over a long period of time.

There is a relationship between the performance of a system and its usability. Usability testing includes testing to verify that the look and feel of the application are good, and that performance of loading of web pages, graphics, etc., is good. There are automated browsing tools which go through all of the links on a page, attempt to load each link, and produce a report including the timing for loading an object or page at various modem speeds. Good usability requires attention to usability in design, and usability engineering is important for web based or GUI applications.

The purpose of post-deployment testing is to ensure that the performance of the web site remains good, and this is generally conducted as part of a service level agreement (SLA). Service level agreements typically include a penalty clause if the availability of the system or its performance falls below defined parameters. Consequently, it is important to identify as early as possible potential performance and availability issues before they become a problem. Thus post-deployment testing will include monitoring of web site availability, performance, security, etc., and taking corrective action as appropriate. Most web sites are operating $24 \times 7 \times 365$ days a year, and there is the potential for major financial loss in the case of an outage of the electronic commerce web site. There is a good account of e-business testing and all the associated issues by Paul Gerrard in detail in [Ger:02].

6.9 E-Commerce Security

The world-wide web extends over computers in virtually every country in the world. It consists of unknown users with un-predictable behavior operating in unknown countries. These users and web sites may be friendly or hostile and the issue of trust arises.

- Is the other person whom they claim to be?
- Can the other person be relied upon to deliver the goods on-payment?
- Can the other person be trusted not to inflict malicious damage?
- Is financial information kept confidential on the server?

Hostility may manifest itself in various acts of destruction. For example, malicious software may attempt to format the hard disk of the local machine and if successful all local data will deleted. Other malicious software may attempt to steal confidential data from the local machine including bank account or credit card details. One particularly destructive virus is the denial of service attack: this is where a web site is overloaded by a malicious attack, and where users are therefore unable to access the web site for an extended period of time.

The display of web pages on the local client machine may involve the downloading of programs from the server, and running the program on the client machine (e.g., Java Applets). Standard HTML allows the static presentation of a web page, whereas many web pages include active content (e.g., Java Applets or Active X). There is a danger that a Trojan horse[4] may be activated during the execution of active content.

Security threats may be from anywhere (e.g., client side, server side, transmission) in an e-commerce environment, and therefore a holistic approach to security is required. Internal and external security measures need to be considered. Internal security is generally implemented via procedures and access privileges.

It is essential that the user is confident in the security provided as otherwise they will be reluctant to pass credit card details over the web for purchases. This has led to technologies such as secure-socket layer (SSL) and secure HTTP (S-HTTP) to ensure security. Another approach is the use of e-cash for electronic payments.

[4] The origin of the term "Trojan Horse" is from the approach used by the Greeks to capture the city of Troy. The Greek hero Odysseus and others hid in a wooden horse while the other Greeks sailed away from Troy. This led the Trojans to believe that the Greeks had abandoned their attack and were returning to their homeland leading behind a farewell gift for the citizens of Troy. The Trojans brought the wooden horse into the city and later that night Odysseus and his companions opened the gates of Troy to the returning Greeks, and the mass slaughter of the citizens of Troy commenced. Troy is located at the mouth of the Dardanelles in Turkey.

6.10 Review Questions

1. Describe the development of the internet.
2. Describe the development of the world-wide web and its key constituents.
3. Describe the applications of the world-wide web.
4. Describe the key constituents of an electronic commerce site.
5. Describe a successful dot com company that you are familiar with. What has made the company successful?
6. Describe a dot com failure that you are familiar with. What caused the company to fail?
7. Discuss the key components of a business model.
8. Discuss software development in a web environment.

6.11 Summary

The invention of the world-wide web by Tim Berners-Lee was a revolutionary milestone in computing. It has transformed the internet from mainly academic use (with limited commercial involvement) to the world of electronic commerce where the use of the internet and world-wide web is an integral part of peoples' lives.

The growth of the world-wide web has been phenomenal since its invention in 1991. New companies were formed to exploit the commercial opportunities of the web, and existing companies developed e-business and e-commerce strategies. Companies providing full e-commerce solutions were concerned with the selling of products or services over the web to either businesses or consumers. The world-wide web is revolutionary in that:

- No single organization controls the web.
- No single computer controls the web.
- Millions of computers are interconnected.
- It is an enormous market place of millions (billions) of users.
- The web is not located in one physical location.
- The web is a space and not a physical thing.

The growth of the world-wide web was exponential, and the boom led to the formation of many "new economy" businesses. The key characteristic of these new businesses was that business was largely conducted over the web as distinct from the "old economy" bricks and mortar companies. Some of these companies were very successful and remain in business. Others were financial disasters due to poor business models, poor management and poor implementation of the new technology.