UNIVERSITY OF LONDON | INTERNATIONAL PROGRAMMES

Data communications and enterprise networking Volume 2

P. Tarr

CO2222

2005

Undergraduate study in
Computing and related programmes

This is an extract from a subject guide for an undergraduate course offered as part of the University of London International Programmes in Computing. Materials for these programmes are developed by academics at Goldsmiths.
For more information, see: www.londoninternational.ac.uk

Goldsmiths
UNIVERSITY OF LONDON

This guide was prepared for the University of London International Programmes by:

P. Tarr

Acknowledgements

I am indebted to David Brownrigg, Ida Pu, Iain Sharpe and my son, Greg, for reading the draft of this subject guide and providing helpful comments and advice. P. Tarr, May 2005

This is one of a series of subject guides published by the University. We regret that due to pressure of work the author is unable to enter into any correspondence relating to, or arising from, the guide. If you have any comments on this subject guide, favourable or unfavourable, please use the form at the back of this guide.

University of London International Programmes
Publications Office
32 Russell Square
London WC1B 5DN
United Kingdom

www.londoninternational.ac.uk

Published by: University of London

Contents

Chapter 1: Introduction

1.1 Aims and objectives

The CIS222 unit aims to provide you with a good grounding in data communications and enterprise networking. It is a Level 2 unit, building on material taught in Level 1 units. The unit concentrates on breadth of understanding rather than depth, and attempts to cover a wide range of networking topics by means of a structured high-level approach. If you desire a deeper understanding of any of the topics covered in this subject guide, you should refer to the recommended texts.

Networks have become increasingly important in recent years. Hardly any applications are now stand-alone. Most applications communicate over a network, either between a client PC and a server in the same building over a Local Area Network or between computers in different cities, countries or continents over a Wide Area Network. The study of networks and applications that make use of networks is now an important part of any undergraduate computer science programme. An understanding of how networks operate will be useful in any computing career you may choose to follow.

The main objective is for you to build a logical framework in which networking topics can be studied, analysed and understood. This will stand you in good stead if you subsequently work in the computer or networking industry and have to design, develop or manage systems that make use of networks. This subject, like many in Computer Science, develops rapidly and new network protocols and technologies are emerging all the time. If you have built a framework which you then can use to analyse and understand these new developments, this will ultimately be of more use than a detailed understanding of current technologies.

A further objective is to help you be able to make informed choices amongst the many competing technologies that are available in the marketplace, if in a subsequent career you are required to make such choices.

The CIS222 unit replaces the CIS208 unit entitled ‘Telecommunications and computer communications’. Some of the material in the CIS208 unit has become very dated, as new technologies have emerged and data transmission speeds have increased way beyond what was imagined possible 10 years ago. Certain promising technologies which were considered important in the CIS208 unit have also failed to come into widespread use, mainly as a result of the phenomenal success of the Internet which did not have a very high profile in CIS208. As a result of this, the new CIS222 unit unlike CIS208 is very much focused on the Internet and its protocols.

1.2 Learning outcomes

On completion of this unit you should be able to:

- describe the technologies deployed in enterprise networks and make informed choices among the competing technologies
- explain how enterprise networks are managed with particular emphasis on performance and security management
- design and cost simple enterprise networks that meet specified requirements
- demonstrate transferable skills in spreadsheet modelling.

1.3 Course outline

The CIS222 unit consists of two distinct parts, which correspond to the two volumes of the subject guide. The first part of the unit is called 'data communications' and is an introduction to the terms, concepts and network architectures required to understand how data is transmitted through communications channels and networks. This, of necessity, requires a technical approach and some knowledge of the physics of transmission as well as the study of network architectures and protocols. The second part of the unit is called 'enterprise networking' and is more concerned with the design and management of networks used by businesses and other large organisations. It therefore concentrates more on the business and management issues that arise.

This volume of the subject guide is more business focused than Volume 1, but does still have a significant technical content. It concentrates on network technologies used by enterprises. It includes a study of the network marketplace and the different products and technologies considered and used by enterprises in designing, building and managing networks. This volume contains eight chapters. This chapter will introduce the unit. Chapter 2 will introduce some marketing concepts and will analyse the main types of company that market network products. The next three chapters will build on material presented in Volume 1 to examine specific technologies used in enterprise networking. Volume 1 followed a top-down layered approach, but in these chapters we will discuss technologies that cross many layers or do not fit neatly into the hybrid reference model that we are using. Chapter 3 will examine technologies used by Personal Area Networks (PANs), Storage Area Networks (SANs) and Local Area Networks. Chapter 4 will examine the technologies used by Metropolitan Area Networks (MANs) and Wide Area Networks (WANs). Chapter 5 will examine how these technologies can be integrated into a single cohesive internetwork and will include studying aspects of IP networks, particularly routing, that were not considered in Volume 1. Chapter 6 will consider the exacting requirements of multimedia networking and the attempts to re-engineer the Internet to meet them. It will also examine Asynchronous Transfer Mode as an alternative technology for multimedia networking. The final two chapters return to a more business-oriented approach. Chapter 7 will consider the task of designing enterprise networks and some of the techniques that network designers use. Finally, in Chapter 8, we will examine how enterprise networks are managed using the ISO[1] Network Management Framework, concentrating on performance and security management, as these are the areas that cause most concern to network managers today.

[1] *International Organization for Standardization.*

This subject guide will contain practical activities that can be carried out to gain further understanding of the topics covered in the unit. These activities can be carried out in a laboratory or at home. You are strongly advised to spend time on these activities, as the extra insight you will obtain will stand you in good stead for the examination, and any coursework assignments.

1.4 How to use this subject guide

This subject guide is intended to provide a logical structure in which to study the subject of enterprise networking. It is not intended to replace your need to read around the subject to improve your understanding. You may wish to follow some of the further reading referenced at the start of each chapter. An alternative or supplement to the further reading is to follow the links on the course web site[2] to study the topics in each chapter.

[2] http://doc.gold.ac.uk/~mas01pt/cis222.

How to best use this subject guide will depend on your personal approach to study and whether you are studying on your own or at an institution.

One suitable approach would be to start with reading a chapter of the subject guide, followed by some of the further reading, if any of the material has not been understood or more information is desired. Then attempt the sample exam questions at the end of the chapter, which can be checked with the model answers and hints in Appendix B, before reading the learning outcomes. If you feel that these have not been fully achieved, go back to the further reading or to the web links, in order to make sure that you have understood each topic. You should also carry out the activities in each chapter, as they provide further insight to the topics being studied.

This subject guide makes use of a large number of acronyms. When a new acronym is introduced, it is expanded in full. Subsequent references will often just use the acronym. A list of acronyms used in the subject guide can be found in Appendix C.

1.5 Reading list

Because of the very varied material presented in this unit, there is no one book that covers the whole unit (or even the majority of it) and that can be recommended for essential reading. You may wish to obtain your own copy of one of the general networking books listed below or use library books and web sites for further reading material.

Canavan, John E. *Fundamentals of Network Security.* (Artech House), first edition, 2001.

Carr and Snyder *Management of Telecommunications.* (McGraw Hill), second edition, 2003.

Fitzgerald and Dennis *Business Data Communications and Networking.* (John Wiley and Sons), eighth edition, 2005.

Forouzan, Behrouz, A. *Data Communications and Networking.* (McGraw Hill), third edition, 2003.

Kurose and Ross *Computer Networking – A Top Down Approach featuring the Internet.* (Addison Wesley),third edition, 2005.

Porter, Michael *Competitive Strategy: Techniques for Analyzing Industries and Competitors.* (Free Press), 2004.

Stallings, William *Data and Computer Communications.* (Prentice Hall International), seventh edition, 2003.

Tanenbaum, Andrew S. *Computer Networks.* (Prentice Hall International), fourth edition, 2002.

1.6 Useful web links

The world wide web is a very dynamic medium and publishing a large number of web links in a subject guide is likely to result in the frustration of a ‘File Not Found’ response at some point in the future. Instead, an up-to-date list of useful links relevant to the unit will be maintained on the CIS222 course web site at

http://doc.gold.ac.uk/~mas01pt/cis222/study/volume_2 .htm

In the event of this URL changing during the life of the subject guide, please follow the links to the author’s home page from the staff page of the Goldsmiths’ Department of Computing web site, and then follow the link to this page.

Two links are particularly useful:

http://www.rfc-editor.org/rfc.html
contains the full text of all the Internet standards which are known as Requests for Comments (RFCs [3]).

[3] Requests for Comments.

http://www.protocols.com/protocols.htm
contains descriptions of most network protocols.

1.7 Study time

If you are studying this unit full time then it should take approximately one quarter of your total study time for an academic year. In a typical institution you will probably spend about four hours per week on each unit during term time in formal teaching. It is recommended that you spend at least three hours per week in reading and private study. You should also on average spend a further three hours per week in attempting sample exam questions at the end of each chapter and in doing the activities, plus formal coursework, in the lab or at home. If you are studying entirely on your own, then you should aim to spend an additional four hours per week in private study, to compensate for the hours that students in institutions receive formal teaching.

You should aim to spend a total of 300 hours on the whole unit, including formal teaching. This total should also include time spent revising during reading weeks, vacations and prior to examinations. This time is applicable to an average student who aims to do well in the unit. Some students who work more slowly may need to devote more time than this.

In order to complete this volume of the subject guide in one (10-week) term, you should aim to complete one chapter per week. This will leave two weeks spare for consolidation at the end or to allow for some slippage.

If you are studying part-time, over a longer period than one term, then you will have to adjust this recommended study time accordingly. Revision for examinations should be in addition to the above.

1.8 Examination

Important: the information and advice given in the following section are based on the examination structure used at the time this guide was written. However, the University can alter the format, style or requirements of an examination paper without notice. Because of this, we strongly advise you to check the rubric/instructions on the paper you actually sit.

The examination will be a single three-hour written paper, usually sat in May, which will consist of a total of six questions. Three questions will be on the first half of the unit (data communications) and three on the second half of the unit (enterprise networking). You will be expected to answer four questions in all: two from the first half of the paper and two from the second half of the paper. Specimen exam questions can be found at the end of each chapter and a complete specimen examination paper can be found in Appendix A. Some model answers and hints can be found in Appendix B.

The overall mark for the unit will be calculated from the examination results and the coursework marks.

Chapter 2: Network markets

Further reading

Carr and Snyder *Management of Telecommunications.* (McGraw Hill), second edition, 2003, Chapter 1.

Porter, Michael *Competitive Strategy: Techniques for Analyzing Industries and Competitors.* (Free Press), 2004, Chapters 1, 2.

This volume of the subject guide is concerned with enterprise networking. An enterprise is simply an organisation. It can be a business (such as a public or private company), a partnership (such as a firm of accountants or management consultants), a not-for-profit organisation (such as a charity or a university) or a government body (such as a government department, a local government council or independent government agency).

Some types of networking are appropriate for both consumer and business (or enterprise) markets. Others are mainly intended for use by enterprises. It is these types of networking that we will be studying in this volume. The difference between enterprise networking and the types of networking associated with the consumer market, is that enterprise networks require more investment in time to design networking solutions and also further effort to manage these networks. The types of networks that support consumers do not require much, if any, design or management by the consumers themselves, but they do require considerable design and management by the network operators and these networks can be thought of as enterprise networks built and managed by network operators to support their customers. These public networks will also be studied in this volume. Enterprise networks can therefore range from that of a small business supporting a few PCs connected to an Ethernet hub, to a very large network owned by a network operator, supporting millions of customers. Enterprise networking is often thought to be limited to data, but enterprises will always have a requirement for voice and video networks, and enterprise networking must also include these, particularly as the integration of these very different types of network is becoming more important.

In this chapter, we shall examine the main drivers that affect the various markets for network products and services. We shall firstly look at this from a marketing perspective and introduce several marketing models that can be used to analyse markets. A study of these markets is useful, as networking professionals not only have to decide on technical solutions, but also have to choose suppliers. It is often the latter as much as the former that determines the best solution. To be able to evaluate suppliers, an appreciation of how markets work and the various strategies employed by the suppliers in the network marketplace is extremely useful.

2.1 Strategies to gain competitive advantage

According to McGaughey, Snyder and Carr, **competitive advantage** is the ability to excel in the marketplace due to price, product, service level or performance. Most enterprises regard obtaining competitive advantage as an imperative. The achievement of competitive advantage ultimately determines whether an enterprise is successful.

Some enterprises, particularly public enterprises, such as government departments or other not-for-profit organisations do not operate in a marketplace and are not driven to seek competitive advantage as there are no competitors. But even these enterprises have constraints on costs that force them to seek improvements in performance, and sensitivity to public opinion which force them to provide services, and service levels that meet exacting requirements. Other public or not-for-profit organisations, such as universities do operate in some sort of competitive market and are driven to seek competitive advantage.

Central to the study of competitive advantage is the concept of a **market**. A market essentially consists of the customers and potential customers to whom a product or service can be sold. A **product** or service is literally anything that is produced, is offered for sale and that satisfies a want or need. In terms of networking, the product usually provides the ability to communicate and is packaged in such a way that it has a well-defined set of functions, service levels and a price (or tariff).

The ultimate size of the market for a product is determined by the strategy of the seller. Companies selling products and services will often restrict their sales to just a part (or segment) of the total potential market. It is this market segmentation strategy that is often crucial to gaining competitive advantage. Before we consider strategies we shall take a closer look at competition.

Michael Porter has identified five forces that drive competition within a market or market segment.

These are:

- the intensity of rivalry among existing competitors – factors that increase rivalry include many competitors, equal market shares, slow growth, high fixed or storage costs, low differentiation, low switching costs and high exit barriers.
- the threat of entry by new competitors – factors that reduce the threat from new entrants include patents, economies of scale, product differentiation, brand loyalty, high capital costs, high switching costs, good distribution channels and government regulation.
- the pressure from substitute products – factors that reduce the threat from substitutes include high switching costs, low prices and brand loyalty.
- the bargaining power of customers – factors increasing the bargaining power of customers include a small customer base, the size of the customer and the size of the orders the customer makes.
- the bargaining power of suppliers – factors increasing the bargaining power of suppliers include the number of potential suppliers, low switching costs and the uniqueness of the product being bought.

All of the above forces will act on a company seeking competitive advantage in a market. The companies must develop strategies to counter these competitive forces.

2.1.1 Generic strategies

Competitive advantage can be obtained using three generic strategies, also identified by Michael Porter. In general, to be successful, the enterprise should choose just one of the strategies below.

Cost leadership

In this strategy, the enterprise seeks to have the lowest costs in its industry, often achieved through economies of scale, operating efficiencies, proprietary technology or preferential access to materials. If cost leadership can be achieved, then the enterprise can offer its products and services to a broad market at the lowest price and thus seeks to achieve or maintain a high market share. It does not attempt to provide unique features and only enhances its products and services in response to market demand.

Low cost does not automatically lead to low price. A cost leader may, if it has a large enough market share, choose to offer its products and services at higher prices and improve its profit margins, but can quickly reduce its prices if threatened by a new entrant or competitor.

There is usually only one cost leader competing within a market. When two enterprises aim to be cost leaders within a market, the result is likely to be a price war from which a single cost leader will emerge.

If an enterprise can obtain cost leadership, it acts as a disincentive to new entrants into the market, as they are unlikely to achieve the same level of costs and if new entrants do arrive, cost leadership can be used to undercut the prices of its competitors.

Cost leadership as a strategy is very vulnerable to new technology. If a competitor can offer a similar product based on a new and cheaper technology, cost leadership and ultimately market share will be lost.

Differentiation

In this strategy, the enterprise attempts to offer products or services, with unique features that customers value, to a broad market. The perceived added value (whether real or not) allows the enterprise to sell its products or services at a premium price. The advantage of this strategy is that the enterprise can maintain its market share because of brand loyalty. This will also deter competitors and new entrants selling substitute products and services.

There can be any number of differentiators competing within a market each with its unique product features.

The disadvantages of this strategy is that competitors may eventually create imitation products and services that may tempt customers away, particularly if there is a price advantage. A differentiation strategy is likely to require continuous innovation. Customer tastes may also change and the value of the differentiating features may diminish over time.

Focus

In this strategy, instead of offering a product or service to a broad market, the enterprise focuses on a particular homogeneous market segment or niche, thus narrowing its market, but aiming to be dominant within its own chosen market niche. The focus within the market segment can either be related to cost or to differentiation. The advantage of focus as a strategy, is that the enterprise can become dominant in the market segment chosen and can maintain a close relationship with its customers and hence better understand their needs and maintain loyalty. The main disadvantage is that the enterprise sells its products and services less widely and thus is subject to higher costs as it is often at the mercy of its own suppliers with whom it has much less leverage than enterprises who target a broader market. The big risk of this strategy is that cost leaders' or differentiators' products or services might also be targeted at the niche market segment.

An initial market segmentation, popular with network operators is to decide whether to aim products and services at individual consumers or other enterprises. Fixed line operators will often call these segments residential and business. New operators often target businesses as they are often more profitable, due to higher usage, geographical location (they tend to be in city centres) and lower selling costs (a single purchaser will often buy on behalf of the whole enterprise).

For consumers the following market segments can be identified:

- **geographic**, based on a particular region, climate, population density or population growth rate
- **demographic**, based on age, gender, social class, occupation, nationality, ethnicity, religion, education, wealth, income or family status
- **psychographic**, based on personality type, hobbies, pastimes, values, attitudes or lifestyles
- **behavioural**, based on usage rate and patterns, price sensitivity, brand loyalty or benefits sought.

For businesses, the following market segments can be identified:

- **geographic**, based on a particular region, customer concentration, regional growth rates and other economic factors.
- **customer type**, based on the type or size of the organisation, its industry or its role.
- **buyer behaviour,** based on organisational structure (centralised or decentralised purchasing), loyalty to suppliers, usage patterns, price sensitivity or order size.

The personal computer market with which you will be familiar, can be used to illustrate the above three generic strategies. In this market Dell, because of its market share and efficient operations that benefit from economies of scale, is able to follow a cost leadership strategy. Many other manufacturers follow a differentiation strategy. Apple Computer is perhaps the best example of a differentiator in this market. It differentiates its products by means of many unique features. Other manufacturers, tend to follow a focus strategy. A good example of a company who now follows this strategy is IBM. Although they invented the de-facto PC standard, they found that they could not compete in the consumer market and pulled out of this market, to focus on the niche business market.

2.1.2 Growth strategies

In order to be competitive and give a good return to its shareholders, an enterprise must seek to grow its business. It can do this within an existing market or create a new market with either an existing product or a new product.

The various possibilities for generating growth were analysed by Ansoff [1] in a matrix.

[1] *H.I. Ansoff 'Strategies for Diversification' in* Harvard Business Review, *1957.*

Table 2.1: Ansoff's Matrix

Product / Market	Existing	New
Existing	market penetration	product development
New	market development	diversification

With **market penetration**, the enterprise markets its existing product, without any alteration, to existing customers by promoting it or by repositioning the brand, with the aim of increasing its share of the market.

With **market development**, the enterprise markets the existing product, without any alteration, to a new market segment. This could be a new geographic or a new demographic market, for instance. With **product development**, the enterprise creates a new product or enhances an existing product and markets it to its existing customers. Finally, with **diversification**, the enterprise markets a completely new product to completely new customers. Diversification can either be related or unrelated to existing products and markets.

These growth strategies have different levels of risk. The order in which they were discussed above is in increasing order of risk.

2.1.3 Product strategies

The Boston Consulting Group created a matrix to analyse the life cycle of products in a company's portfolio, in the form of a matrix, known as the Boston Matrix.

Table 2.2: The Boston Matrix

Market Share / Market Growth	High	Low
High	Star ←	Problem child
Low	↓ Cash Cow →	Dog

The matrix compares market growth and market share and can be used to track a product's desired life cycle, shown by the arrows. The terms it coined, such as cash cow, are now very well known.

When an innovative new product is launched, it will often have a low market share but hopefully high market growth. If so, it is a **problem child**, because its launch will have been costly and the revenue obtained from its sale will not recover these costs for a long time. If the product does not have high growth, then it is unlikely to be profitable and hence it is a **dog**. If it does have high growth, it will obtain a high market share and hence it is a **star**. Stars are profitable, but maintaining growth is also costly. But eventually, there are limits to how much growth can take place and the rate of growth must at some time slow down. At this stage with a high market share and low growth, it is a **cash cow**, so named because the revenue from the product can be 'milked' without much effort. Finally, the product is likely to lose its market share to substitute products at which point it becomes a dog.

The ideal strategy is to use this matrix to analyse the position of all the products in a company's portfolio of products and to plan the launch of new products in such a way that the portfolio contains a balance of each type (except dogs). This means that the revenues from cash cows can be used to develop new problem children and turn them into stars and then cash cows to replace the existing cash cows when they become dogs. A company cannot rely on milking its cash cows for ever. These will ultimately turn into dogs as competitors offer substitute products and steal market share. Companies must continually innovate new products to survive.

2.1.4 Partnership strategies

In order to obtain or retain competitive advantage, it is often necessary to form partnerships or alliances with other companies within the same or related markets. Participation in forums to develop standards is one example of this. Without clear standards, customers are often reluctant to buy products due to fear that the market will move in a different direction and they will end up with a product that few others are using. This is particularly true in the area of networking, where customers want to be able to mix and match products from different suppliers, so that they are not locked into a single supplier and thus can gain bargaining power over their suppliers.

Partnerships are often required to develop innovative products. Research and development is an expensive activity and by collaborating with partners in developing new products both the cost and the risks can be shared.

Companies also sometimes enter into partnerships with their suppliers or even their customers. Such partnerships are often necessary to distribute products, if the company which has developed the product does not have the geographical reach or organisational structure in place to distribute the product.

2.2 Market studies

In this section, we shall firstly introduce some tools (PEST and SWOT analysis) that can be used to analyse an enterprise and its business environment. We shall then go on to study the fixed and mobile network markets, followed by the ISP[2] market and the network equipment market. Finally, we shall have a look at the market for a network application in an industry sector, the Global Distribution Systems (GDSs), used by the travel industry. This has been chosen, as it is one of the few industry sectors that has built a large global network that you may have used, unknowingly, when booking holidays.

[2] Internet Service Provider

PEST analysis examines a companies environment under the following headings:

- **political** factors such as government policies and the regulatory environment (which is particularly important for telecommunications services)
- **economic** factors such as interest rates, inflation rates, unemployment rates, disposable incomes and economic growth rates
- **socio-cultural** factors such as customer attitudes, leisure time and demographics
- **technological** factors such as technical innovations, patents and the development of new systems and distribution channels.

SWOT analysis is a subjective examination of a company's Strengths, Weaknesses, Opportunities and Threats. PEST analysis and Porter's five forces listed in section 2.1 and the company's generic, growth, product and partnership strategies discussed in Sections 2.1–2.4 can be used to help identify some of the factors under each heading.

- **Strength** are positive aspects internal to the organisation.
- **Weaknesses** are negative aspects internal to the organisation.
- **Opportunities** are positive aspects external to the organisation.
- **Threats** are negative aspects external to the organisation.

You should use these tools to analyse companies that operate within your country in the activities at the end of each section. The subject guide attempts to describe the various network markets from a global rather than national perspective.

2.2.1 Fixed network operators

Fixed network operators are usually classified according to whether they are incumbents, who historically provided telecommunications services in a protected market or insurgents who entered the market when it became open. Thirty years ago, it was commonly believed that telecommunications services within a country were best provided by a single operator, as telecommunications along with other utilities were regarded as natural monopolies. These operators were often owned by governments or else were private companies that governments entrusted to run all the telecommunications services in their countries. The PTTs[3] and RPOAs[4] were all powerful within their countries and would often insist that customer premises equipment (CPE) such as telephones, PABXs[5] and modems, were provided by themselves. Some would allow customers to buy their own CPE and connect it to the network, but only after the equipment had undergone rigorous 'permission to connect' testing which aimed to ensure that the equipment was totally safe and would not apply a mains voltage to the line and electrocute one of the PTT's employees. Because of these and other monopolistic and restrictive practices there was no competitive market in telecommunications. PTTs were slow to innovate and tended to provide a restricted number of low quality services at excessively high prices. The profits made by the PTTs were often regarded as an extra source of revenue by governments and the PTTs found it hard to obtain capital for investment in new technology, as it would increase government borrowing. Many PTTs did not have sufficient line plant to meet demand and they imposed a waiting list for service and line sharing schemes.

[3] The Government Departments that ran posts and telecommunications as monopolies were known as Post, Telegraphs and Telephones (PTTs).

[4] Private companies that ran telecommunications services as monopolies were known as Recognised Private Operating Agencies (RPOAs).

[5] Private Automatic Branch Exchange – a customer owned and managed telephone switch.

There were however some positive aspects to this situation. The PTTs were engineering led and as a result constructed and designed networks that were technically sound, within the constraints of the technology of the times. The markets and prices were stable and growth was steady, if unspectacular. The organisations often had a good public service ethos, as befits a part of government and they were obliged to provide service in any part of their countries. This is known as a **Universal Service Obligation (USO)**. Finally, by working together in CCITT[6], which became the ITU-T[7], they were able to develop inter-working standards. The fact that the global telephone network is probably the largest single system mankind has ever built and that all of the individually managed networks can interconnect and inter-work with each other is a testament to this standardisation process.

[6] Comité Consultatif International Téléphonique et Télégraphique.

[7] International Telecommunications Union – Telecommunications Standardization Sector

The operators forged bilateral interconnection agreements with each other, which required the operator on whose network a call originated to pay the operator on whose network a call terminated a per minute charge. This charge was not related to cost and for many years the operators made huge profits out of international telephone traffic paid for by their captive customers, whom they usually called subscribers, thus reflecting their attitudes to customer service.

All this began to change in the 1968, when the US Federal Communications Commission allowed other businesses to connect their equipment to the **incumbent** (monopoly) network owned by AT&T[8]. This decision was shortly followed by another ground-breaking decision that allowed a small upstart company called MCI[9] to provide long distance telecommunications services

[8] American Telephone & Telegraph who were the RPOA that owned the Bell network in the USA.

[9] Microwave Communications Incorporated.

in competition to AT&T. MCI started out by setting up a high speed communications network using microwave links between major US cities, but had to fight AT&T all the way through the courts and by lobbying Congress, until eventually won an important anti-trust case in 1984 which resulted in AT&T reorganising its 22 Bell Operating Companies and divesting them into seven separate Regional Bell Operating Companies (RBOCs)[10], and a long-distance (inter-state) carrier which retained the name AT&T. At the same time AT&T also divested itself of Bell Labs and its manufacturing arm Western Electric[11]. MCI and other new carriers such as GTE, Sprint and Worldcom were then able to compete fairly with AT&T for the long distance market. But within each region there the RBOC still had an effective monopoly. The RBOCs, who were also known as Incumbent Local Exchange Carriers (ILECs), gradually merged with each other and other new operators to form the US operators that we know today such as BellSouth, Qwest, SBC and Verizon. At the time of writing both Qwest and Verizon are bidding to take over MCI and SBC has bid to take over AT&T. Other companies known as Competitive Local Exchange Carriers (CLECs) were allowed to compete with the RBOCs for local lines and calls.

[10] Also sometimes referred to as 'Baby Bells'.

[11] Now Lucent Technologies.

Meanwhile in Europe, things were beginning to change in the UK. The British Government had split the Post Office's telecommunications business from its postal business and began to privatise it as British Telecommunications (BT) plc[12]. At the same time just one new operator called Mercury Communications Ltd (MCL) was licensed to compete with BT. This was owned by a consortium which included Cable & Wireless (C&W) Limited, which had also only recently been privatised and was the incumbent operator in a number of British colonies and ex-colonies (notably Hong Kong). MCL built a modern fibre optic network alongside railway tracks and attempted to compete against BT in all its markets, including unprofitable pay phones. As a result, this strategy was not particularly successful. MCL should have focussed on profitable niches rather than attempting to compete in all markets. MCL was eventually forced to focus on niche business markets and the residential parts of MCL's business were eventually sold off to the cable TV company NTL with the business parts being re-absorbed into C&W.

[12] Now BT Group plc.

In 1991, the British government abandoned its duopoly policy and announced that there would be no restrictions on the number of telecommunications operators that would be licensed. This led to a large number of new entrants to the UK market which eventually consolidated into a smaller number of successful operators such as Colt (which started out as City of London Telecommunications), Energis (which ran its fibre optic cables over the UK electricity grid) and Thus (which started out as Scottish Telecom).

The other players in the UK fixed telecommunications market are the cable TV companies. These were started up in the 1980s, with heavy investment from the US RBOCs, to provide cable TV and telecommunications services in franchised areas. The heavy costs of laying cables in residential areas and competition from satellite TV has caused many problems for the cable TV companies in the UK and they have now consolidated into NTL and Telewest and may consolidate further.

The UK policy of privatising its PTT has now been followed by most countries throughout the world (sometimes with partial privatisation) and has changed the face of the telecommunications market. It is generally felt that competition leads to lower prices, more innovation and better service. Also, governments, as the owners of the PTTs, were able to collect large revenues when the newly formed telecommunication companies were floated. Many of the well-known global players have their origin as PTTs. These include BT in the UK, Deutsche Telekom in Germany, France Telecom in France, NTT in Japan, Telecom Italia in Italy, Telefonica in Spain, TeliaSonera in Scandinavia, Telstra in Australia and SingTel in Singapore. Many countries however have not followed the UK and US policy of encouraging infrastructure competition where new companies are licensed and encouraged to build new large-scale local networks. Instead the emphasis in many countries has been on service competition where the incumbent is heavily regulated and has to offer wholesale services to service providers who then sell to customers in competition with the incumbent itself. As a result of this policy, in many countries, there is only a limited amount of infrastructure that is not owned by the incumbent, but providing regulation is effective, this does not prevent a vibrant competitive telecommunications services market developing. Most regulators have ensured that all operators are given equal access to the incumbent's local loops.[13] This has been achieved primarily through forcing the incumbents to offer carrier pre-selection (where customers can pre-select their local, national and international carriers without having to dial access codes), Local Loop Unbundling (LLU). With LLU, the incumbent's local telephone switch is bypassed altogether and the loop is terminated on equipment belonging to another operator) and number portability (where customers can retain their numbers when they switch between operators).

[13] The pair of cables that connect the customer's premises to the local exchange.

During the 1980s and 1990s there was a boom in the telecommunications industry and many new entrants emerged intent on building new infrastructure. As a result new operators, such as Interoute, Global Crossing and Level 3, have laid optical fibres in many parts of the world and with the ever increasing transmission speeds and bandwidth which can be used in the existing fibres, there is now a glut in capacity and as a result bandwidth prices have fallen and network bandwidth has become a commodity product. This has been something of a problem for the new operators.

Amongst network operators, the incumbent (ex monopoly) operator is rarely the cost leader. This is because they tend to have large overheads and they own an expensive inefficient legacy network, as well as having a Universal Service Obligation (USO) which means that they have to provide services to parts of the country that are uneconomic. The strategy of incumbents is usually to attempt to maintain their large market share by differentiating themselves on service with slightly higher prices than the new entrants. The incumbents have two main cash cows, resulting from their legacy investment in infrastructure. These are voice calls over the PSTN[14] and the sale of private circuits. Both of these are mature markets which have limited growth opportunities. In developed countries, the number of customers for the PSTN voice has reached the point of saturation, where nearly everyone has a phone line. Because of this lack of growth the fixed network operators are looking for new stars and see broadband access in this light. It has high growth potential. They also are looking to diversify, particularly with business customers, where they aim to provide systems integration and outsourcing services adding value to basic telecommunications or to offer new mobile or IP based services. In many countries, the fixed network operators have used their market strength, customer base and distribution channels to successfully diversify into these related markets and many are now the dominant mobile network operator and Internet Service Provider in their countries. Eventually, both PSTN voice and private circuits are likely to become dogs as voice calls are substituted by Voice over IP (VoIP) and to some extent mobile phone calls and e-mail. Private circuits are also likely to be substituted by broadband access to managed services. This presents an interesting dilemma for fixed network operators as they have to decide at

[14] Public Switched Telephone Network, see Section 4.2.4 of this volume.

what point to start abandoning their cash cows and launching services such as VoIP (Voice over IP) which will compete against their own traditional PSTN services. This process of companies launching products that compete against their own cash cow products is called **cannibalisation**. If they do not do this, they risk losing everything, as competitors will certainly launch these services. A similar dilemma exists between traditional private circuits and ATM[15] or IP services provided over broadband access.

[15] Asynchronous Transfer Mode, see Section 6.5 of this volume.

The new entrants often known as **insurgents** or alternative network operators (or **altnets**) have had to focus on niche markets (usually in the business sector and in restricted geographical areas such as city centres) in order to survive. They are often much more innovative and adaptable than the incumbents. They have the advantage of not being encumbered by legacy technology, regulation or a universal service obligation. Their lower cost base, resulting from this, means that they have the potential to be cost leaders and to implement a cost focus strategy to compete against the incumbent on price, although start-up costs, particularly investment in infrastructure act as barriers to entry into the market, as do regulatory hurdles. Unlike the incumbents, the new entrants are able to achieve high growth for their PSTN voice and private circuit offerings, as they encourage customers to switch from the incumbents by their lower prices. These services, as well as newer services such as broadband, are all stars for the new entrants.

Many fixed network operators have tended to restrict their operations to their own country, but have involved themselves in partnerships with each other to develop new standards, launch satellites and form bilateral agreements for international services.

Those that have strayed outside their own countries, have often done so in partnership with other operators. Examples of these partnerships are Concert (now split up between AT&T and BT), Global One (now merged with Equant and owned by France Telecom), Infonet (being acquired by BT at the time of writing) and Unisource (whose network assets were acquired by Infonet). However most of these partnerships ultimately failed and reverted to single ownership, mainly due to problems between the partners, often over the control of the partners' home customers who also had global interests.

In summary, competition in fixed telecommunications markets has certainly led to large reductions in prices, more choice for customers, more innovation and improved levels of customer satisfaction.

Activity 2.1

Make a list of the main fixed network operators who provide services in your country. Look at their web sites and other information on the web to discover what markets they are focused on, what products they are selling and what strategies they are following. Carry out Five Forces, PEST and SWOT analyses on each of the companies.

2.2.2 Mobile network operators

The mobile network operators have only existed since the 1980s. The earliest mobile phone networks were built in Scandinavia where the costs of providing fixed communications to remote communities were prohibitive. The mobile network operators tend to be young dynamic companies that have created a whole new market and have seen it grow at a phenomenal rate. Some mobile network operators are owned by fixed network operators, but where this has happened, the mobile operation is usually run at arms length from the fixed line business. Often this is for regulatory reasons, as regulators want to ensure that the operators do not cross-subsidise their mobile operations from their monopoly revenues. But also, most fixed line operators realise that the culture of a mobile operator is very different and that the mobile operation needs to be much more dynamic and entrepreneurial, which will be much harder to achieve within an ex-PTT organisation.

Another difference between mobile and fixed network operators is that with the former there is usually no incumbent, because governments have, in most cases, issued multiple licenses to encourage competition. All operators have had to start with a clean sheet and build their networks and their customer base from scratch. Building a mobile network requires a huge investment in infrastructure (radio masts, base stations, mobile switches and private circuits which have to be leased from fixed network operators). In addition to this, governments have recently realised that they can generate huge revenues from the sale of radio bandwidth and the licenses to use a part of the radio spectrum are now often auctioned, so that mobile network operators have to compete for them. Despite the large investments, the rewards are also great. Customers are willing to pay a premium price for the flexibility to make mobile phone calls. There is also a revenue stream that comes from calls made between fixed networks and the mobile networks. The fixed network operator who bills the call, has to pay a termination charge to the mobile network operator who receives the call and vice versa. There are further opportunities to make revenue from roaming, where the customer makes or receives calls on a mobile network in another country. Another source of revenue comes from text messages and the downloading of ring tones. Finally, with the advent of data services and third generation mobile phone services, the mobile network operators are also expecting to make significant revenues from always-on data services.

Because there is no incumbent and competition is both fierce and fair, mobile network operators tend to use a differentiation strategy. They attempt to launch new features and services that will give them an edge (albeit a small one) over their competitors. They are also notorious for developing complicated tariff structures that make it very difficult to compare prices between operators. A common technique for subscription accounts is to provide handsets free or at a reduced price, but with customers paying a higher monthly subscription. This was vital, in the early days of mobile telephony, when handsets were very expensive and customers were put off by the up-front cost of buying a handset. An important and fairly recent innovation that has led to increased market penetration, particularly amongst the younger generation who do not have the credit history to open a subscription account, is the use of pre-payment where there is no monthly subscription, but calls are paid for in advance by purchasing a card. The use of pre-paid cards has increased the market for mobile phones, which was once for businessmen and the wealthy professional classes, so that now it is a truly mass market which covers the whole population.

To a mobile network operator, voice services are all stars, possibly to become cash cows in the near future. There is still a good potential for growth left in the market as they can continue to grow their subscriber bases as well as growing revenue per user. There are signs however that growth in subscribers in some countries is tailing off, as a saturation point is close to being reached. Future high growth is dependent on the success of the related diversification into new data and third generation services. These services are currently problem children with a high growth rate and the potential to become stars.

Because the mobile phone market is so competitive and there is little difference between the product offerings of the operators, there is a tendency among customers to switch operators. The movement of customers from one operator to another is known as **churn**. Operators try to prevent churn by locking customers into lengthy contracts and by selling handsets that can be blocked from use on other operators' networks.

Mobile network operators have historically tended not to form partnerships with each other in their own markets, but they have entered in partnerships with suppliers and distributors and with other international operators for roaming. There has also been considerable international takeover activity between mobile network operators (sometimes hostile) and expansion into new countries, forming large multinational operators such as Orange (owned by France Telecom), T-Mobile (owned by Deutsche Telekom), Hutchison Whampoa Limited (based in Hong Kong, but successfully launched Orange in the UK before selling it to France Telecom and acquiring several 3G licenses around the world, marketing under the global brand name '3'), NTT DoCoMo in Japan, who have been highly successful in launching I-mode data services, and Vodafone, which was just a start-up company in 1985 when it won one of the first analogue licenses in the UK. Vodafone has been most successful in its expansion and acquisitions, and after acquiring Airtouch in the US, it is now the global market leader. It is the largest company listed on the London Stock Exchange with a market capitalisation many times larger than BT (the UK's largest fixed network operator).

The mobile network operators have therefore consolidated into a small number of global brands, in contrast to the fixed network operators who were predicted to do this, but have yet to achieve it.

Activity 2.2

Make a list of the main mobile network operators who provide services in your country. Look at their web sites and other information on the web to discover what markets they are focused on, what products they are selling and what strategies they are following. Carry out Five Forces, PEST and SWOT analyses on each of the companies.

2.2.3 Internet Service Providers

The Internet Service Provider (ISP) market was created in 1994, when the US Government commercialised the Internet. Up until then, the Internet's NSFNET[16] backbone could only be used for educational and research purposes. Commercial providers were encouraged to build or acquire and operate their own backbones and regional access networks, to interconnect with each other via Commercial Internet Exchanges (CIXs), and then to sell their services to the general public and the business community. Some of the companies that became involved in this were fixed network operators, such as Ameritech (acquired by SBC), MCI, MFS (acquired by WorldCom), Pacific Bell (also acquired by SBC), Sprint, WorldCom (who acquired MCI and eventually reverted to the MCI brand). In other parts of the world the fixed network operators, such as France Telecom (with Wanadoo), Deutsche Telekom (with T-Online) also provided ISP services (usually via a subsidiary which can be dynamic and entrepreneurial), and in many countries they have become the dominant ISP, being able to exploit the large customer bases of their parents. A number of ISPs were start-up companies, such as Performance Systems International (PSI) in the US and Demon Internet (acquired by Thus), Pipex (now owned by MCI) and FreeServe (acquired by Wanadoo) in the UK, who foresaw the commercial potential of providing subscription-free internet access. Others providers were the on-line

[16] *National Science Foundation Network.*

computer service/information providers, such as America On-Line (AOL), CompuServe (acquired by AOL), Prodigy and UUNet (also now owned by MCI). AOL has been able to build a very large global customer base through aggressive marketing and international expansion.

The ISP market is highly competitive. The products it offers are in great demand and growth has been spectacular, but still with potential to grow further, as market penetration improves. They are certainly stars and ISPs have had to differentiate themselves by means of continuous innovation. One of the innovations, which has made a big difference to the market was the provision of 'free' internet services. In the early days, in most countries, customers paid their ISP a monthly subscription and then paid their fixed network operator a per minute telephone call charge[17] whenever they dialled-up their ISP. In the UK a start-up company called FreeServe (now owned by and rebranded as Wanadoo) realised that they could offer a subscription-free service by offering Internet access via a local call fee access number where revenue from the billed call was passed onto the ISP by the fixed network operator. This was an attractive proposition for customers, who could now access the Internet without having any monthly paid contract with the ISP. As a result of this and using a good retail distribution channel, FreeServe was able to come from nowhere to be the market leader in the UK (for a period). Tiscali enjoyed similar success in Italy and has now expanded, mainly by mergers and acquisitions to cover most of Europe.

[17] *In North America, and now some other countries too, local telephone calls are not charged on a per minute basis. The cost of calls is recovered through a higher monthly rental charge and as a result ISP customers are not charged for connect time and can afford to remain on-line for long periods of time.*

Since then, the ISP market has largely reverted back to a subscription based model, as there was huge pressure from customers to remove per minute phone call charging which was a disincentive to use the Internet. This pressure built up on the fixed network operators who were forced to offer un-metered dial-up Internet access via wholesale products for the ISP to offer to their customers. The same model is also used for broadband Internet access, although some products restrict the volume of data that can be downloaded.

ISPs are also beginning to diversify into offering services that have traditionally been offered by fixed network operators. Such services include broadband access taking advantage of Local Loop Unbundling, or wholesale products. The ISPs are often more innovative than the incumbent operator (or its ISP subsidiary) and can differentiate themselves with higher speed connections and new features. They also compete aggressively on price. They are beginning to offer a VoIP service that will be able to substitute normal telephone calls which may ultimately turn the incumbent's cash cows into dogs.

In summary, the ISP market is still a young dynamic market experiencing high growth rates which appears to have potential to compete with the incumbents' telephone services.

Activity 2.3

Make a list of the main ISPs which provide services in your country. Look at their web sites and other information on the web to discover what markets they are focused on, what products they are selling and what strategies they are following. Carry out Five Forces, PEST and SWOT analyses on each of the companies.

2.2.4 Network equipment manufacturers

The other big players in the network market are the equipment manufacturers. A number of these such as Alcatel, Ericsson, Fujitsu, Lucent Technologies, Marconi, NEC, Nortel and Siemens, began as telephony equipment manufacturers selling to the PTTs and network operators. Many of these companies were originally based in just one country and had a cosy relationship with their incumbent operator from whom they milked a cash cow. However, with increasing competition in these markets, both new entrants and incumbents were prepared to buy from alternative foreign suppliers and the network equipment market was globalised. The countries of origin of these now global companies are still apparent in the location of their head offices. These companies have tended to continue to sell WAN equipment to network operators and other companies, although some of them, notably Ericssons and Siemens, have carried out a related diversification into mobile network equipment. One company, Nokia has re-invented itself several times in its history. Sometimes the diversification has been related, sometimes it has been unrelated. Nokia was founded in 1865 as a paper manufacturer. It diversified into a conglomerate company that manufactured many different products, including rubber, chemicals and telecommunication cables. It later carried out a related diversification into manufacturing telecommunications equipment. In 1992, Nokia sold all its other businesses and focused entirely on telecommunications. It was in the right place (Scandinavia) at the right time (when mobile phones were about to be launched) and it was able to capture a large share of both the mobile handset and the mobile base station and switch markets. As a result of this it is the dominant supplier of mobile handsets in the world.

Several other manufacturers were founded more recently to address gaps or create products in the markets for LAN or internetworking equipment. Such companies include 3Com, Cisco Systems, Foundry Networks and Juniper Networks. Of these companies 3Com and Foundry Networks started out in LAN hub and switch markets, but are moving into in routers and layer three switches, while Cisco Systems and Juniper Networks started in the router market (Cisco invented the router and is still dominant in this market), but has diversified into LAN switches.

Cisco Systems

Cisco Systems is an extremely successful company. It has become the leader or the second player in every market in which it competes. It has achieved this mainly by means of clever acquisitions. Cisco is always on the look-out for small start-up companies with good ideas for promising new products. It then acquires them, integrates them and continues to develop the new products finally marketing them under the Cisco brand. Cisco has carried out this process many times and has a lot of experience at integrating these start-up companies. Many of their executives have stayed with Cisco, continuing to develop their product and some have also become Cisco executives. Cisco has also focused on excellent customer service and has practised what it preaches in terms of using networks to gain competitive advantage. Cisco has developed its business systems to be accessible over the web and employees, customers and suppliers are able to do most of their business with the company using these systems. Cisco always has a portfolio of products at different stages of their product life cycles where the cash cows can provide the revenues to fund the problem children and the stars. Cisco also successfully outsources some of its activities such as manufacturing, but is able control its suppliers closely and ensure timely delivery and quality. Cisco also has put together a certification programme for engineers who design or maintain Cisco networks. This programme has become an industry standard networking qualification and is also taught at many educational institutions. It gives Cisco a real competitive advantage, as anyone studying for these qualifications becomes familiar with Cisco's products. It has used its certification programme and technical forums on its web site to create a community of experts who will often support each other without involving Cisco's own technical support.

Juniper Networks

Juniper is a start-up company, founded in 1996, that has focused on the market for high specification routers which it develops and sells to ISPs and network operators. No company can tackle Cisco head-on in all its markets, but there is room for competition in niche markets such as this.

Alliances

As a result of intense global competition there has been a great deal of acquisitions and mergers amongst network equipment manufacturers and the market has consolidated so that there is now a relatively small number of global suppliers who offer a wide range of networking products. These suppliers compete aggressively with each other, but they also have to collaborate on the development of standards and sometimes on new products which are expensive to develop. Most of these suppliers participate in standards bodies[18] such as the IEEE, the IETF and the ITU-T as well as various forums to progress particular standards such as the ATM and Frame Relay Forums. The reason they do this is because customers are generally reluctant to buy networking products until standards are agreed and suppliers do not want to develop products that are then made obsolete by standards. By participating in standards bodies the suppliers aim to influence the standards in ways that may favour their companies, as well as assisting in developing markets for new products.

[18] See Volume 1, Section 3.3.

Activity 2.4

Explore Cisco Systems' web site (www.cisco.com) and investigate the factors that have given Cisco Systems competitive advantage in its chosen markets.

2.2.5 Global Distribution Systems

Network-based Computerised Reservation Systems (CRSs), or Global Distribution Systems (GDSs) as they are now called, were originally developed for internal use by the airlines in the 1960s. American Airlines came up with the idea of providing terminals to travel agents and a network to connect them to their mainframe computer and started to roll out this network in the 1970s. It called this system SABRE (which originally stood for Semi-Automated Business Research Environment), which it had jointly developed with IBM. This innovative idea had many advantages for American. The main advantage being that travel agents could make the bookings direct without American having to employ staff in call centres to answer telephone calls from the agents and make the bookings on their behalf. Secondly, by providing the terminals and the network, it encouraged loyalty from the travel agents who would often choose American flights for their customers in preference to other airlines.

As far as the agents were concerned the CRS network had arrived at just the right time. The US airline industry was about to be deregulated. Many new airlines came into existence, competition became cut-throat, prices were changing rapidly. Passenger numbers were expanding and the major airlines were moving towards a hub and spoke route structure, which meant that more booking for connecting flights were required. The existing manual ticketing systems and procedures just could not cope much longer, so the travel agents were desperate to connect to American's CRS network, as it was the only one available. This gave American a huge competitive advantage over other airlines who were at least a year behind in deploying similar networks. Given this head start with a new product in the new CRS market, American were soon able to turn their SABRE from a problem child into a star. Having achieved market dominance, SABRE diversified to support hotel, car-hire and other airline bookings in addition to its own flight bookings. It was able to demand a fee from any other company whose bookings could be made using SABRE as it owned the main distribution channel to the travel agents. As a result of this, SABRE was a huge commercial success and soon American was earning more profit from SABRE than it was from its flights. It created a separate company for this new business, which reported along with American Airlines to a holding company called AMR. In 1992, Robert L. Crandall, the CEO of AMR said: 'If you told me I had to sell either the airline or the system, I'd probably sell the airline'.

The other airlines, who were at a considerable disadvantage, decided to develop their own systems in partnerships (United and Eastern collaborated with IBM to produce a system called PARS which further developed into three separate systems: Apollo used by United; Worldspan used by Delta, Northwest and TWA; and SystemOne used by Continental and Eastern and later Texas Air).

With market dominance, American was able to exploit SABRE to the advantage of its airline. When displaying possible flights to travel agents, SABRE would present American flight at the top of the screen and competitors' flights further down. Agents were thus more likely to book flights with American. They also charged airline competitors more than partners for holding information on SABRE. This was clearly anti-competitive and eventually American's competitors won an anti-trust case that prevented American from continuing these practices. This created a more competitive market.

In Europe and Asia Pacific, American were not able to obtain market dominance, as the airlines in these parts of the world developed their own systems based on those already developed by the American carriers. A consortium of European airlines including BA, KLM and Aer Lingus developed Galileo based on Apollo and another consortium including Air France, Iberian and Lufthansa developed Amadeus based on SystemOne. In the Asia Pacific region, Singapore Airlines and Cathay Pacific and others developed Abacus based on SABRE.

The Global Distribution Systems, as the CRSs became known, were probably one of the first examples of e-commerce, even though they started with mainframe transaction processing technology, low speed modem networks and the limited block graphics of dumb teletext terminals. They now exploit web technologies, high speed networks and high resolution graphics supported by PCs.

Specimen examination question

(a) State whether each of the following statements is true or false and, if false, correct the statement:

i. A problem child has low market growth and low market share

ii. Unlike the mobile network operators, the fixed network operators have not been consolidated into global companies.

iii. It is virtually impossible for the incumbent fixed network operator to become the dominant ISP in its home country because it lacks the necessary dynamism.

v. Network equipment manufacturers usually work together in standards bodies and industry forums as customers will often not buy products until standards have been agreed.

(b) List the five forces that drive competition within a market as identified by Michael Porter.

(c) Describe the main difference between the markets of the fixed and the mobile network operators.

(d) Describe the factors that enabled Nokia to become dominant in the mobile handset market.

(e) Describe the factors that gave American Airlines competitive advantage when they connected travel agents to their SABRE network.

Learning outcomes

At the end of this chapter, you should be able to:

- describe the five forces that drive competition in a market, as identified by Michael Porter
- describe the various strategies that companies employ to gain competitive advantage
- analyse network markets using SWOT analysis combined with PEST analysis, Porter's five competitive forces and the strategies used to gain competitive advantage
- describe the historic development and strategies employed by the fixed network operators
- describe the historic development and strategies employed by the mobile network operators
- describe the historic development and strategies employed by the ISPs
- describe the historic development and strategies employed by the network equipment manufacturers
- describe the historic development and strategies employed by the Global Distribution Systems.

Notes

Chapter 3: Network technologies – PANs, SANs and LANs

Further reading

Fitzgerald and Dennis *Business Data Communications and Networking.* (John Wiley & Sons), eighth edition, 2005. Chapters 6, 7, 8.

Forouzan, Behrouz A. *Data Communications and Networking.* (McGraw Hill), third edition, 2003. Chapters 14, 15.1, 15.2, 16.3.

Stallings, William *Data and Computer Communications.* (Prentice Hall International), seventh edition, 2003. Chapters 15.3, 16.2, 16.4, 17.2.

Tanenbaum, Andrew S. *Computer Networks.* (Prentice Hall International), fourth edition, 2002. Chapters 4.3, 4.4, 4.6, 4.7.6.

In this chapter and the next chapter, we will examine the technologies used in enterprise networking. Some technologies will also be discussed that are predominantly used in consumer markets, but they are also used by employees of enterprises for business purposes and it is appropriate to consider them in this part of the course. Many of these technologies cover more than one layer of the hybrid reference model, so they do not fit neatly into the structure of Volume 1 of the subject guide, which considers each layer in a separate chapter.

For convenience, we will classify the technologies by the type of network that they are designed to support. The types of network are Personal Area Networks (PANs), Storage Area Networks (SANs), Local Area Networks (LANs), Metropolitan Area Networks (MANs) and Wide Area Networks (WANs). All of these, apart from SANs were introduced in Volume 1 of this subject guide. The type of network is characterised by the range of distances over which they operate. Table 3.1 below shows roughly how the different types of network compare with each other against a logarithmic scale of distances.

Table 3.1: Approximate ranges of each type of network

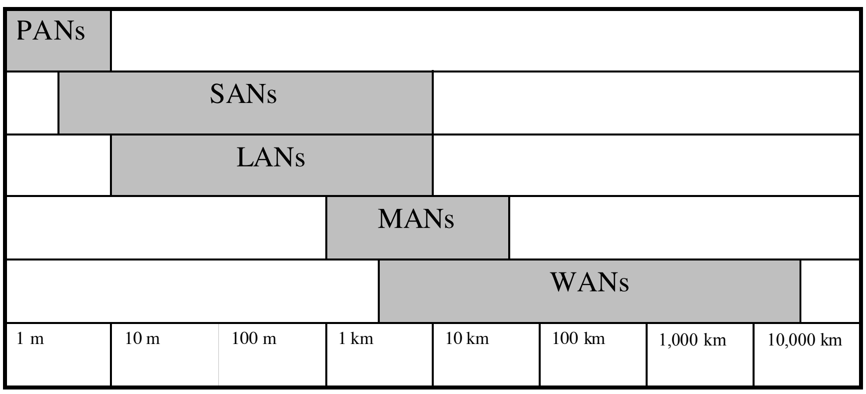

In this chapter we will look at PANs and consider the Universal Serial Bus (USB) and FireWire as examples of wired PANs and Infra-red Data Association (IrDA), Bluetooth and IEEE 802.15 as examples of wireless PANs.

We will consider the special networking needs of SANs and Fibre Channel as the main technology used to support these networks.

Finally, we shall examine the main technologies used in Local Area Networks, concentrating on the many variants of Ethernet, but also considering Token Ring and Fibre Distributed Data Interface (FDDI).

3.1 Personal Area Networks (PANs)

Personal Area Networks (PANs) typically cover a short range (a few metres) close to a person or their personal equipment such as PCs, Personal Digital Assistants (PDAs) or telephones. PANs can be wired or wireless. Wired PANs include external computer buses such as the Universal Serial Bus (USB) or Firewire. Wireless PANs include infra-red communications as sometimes used between PCs and keyboards as well as high-speed radio communications used between intelligent mobile devices.

3.1.1 Universal Serial Bus (USB)

USB has a maximum data rate of 12 Mbit/s (480 Mbit/s for USB 2) and can support up to 127 devices on the same external bus, which consists of up to five metres (without repeaters) of four pairs of UTP[1] cable. It can support multiple devices at 1.5 Mbit/s or one device at 12 Mbit/s. Multiple devices can be daisy-chained by being connected to other USB devices or daisy-chained from a USB hub in a bus-star topology. USB also distributes low voltage power to peripherals, which no longer require separate low voltage power supplies. USB supports plug and play attachment (i.e. automatic assignment of resources to devices) and all devices are hot pluggable (i.e. that can be attached and start working without crashing or having to restart the PC). When a new device is plugged into the bus, the change in power levels is detected and an address is assigned to the device, which is unique on the bus. USB avoids contention by operating a master–slave protocol. USB uses NRZI[2] encoding with bit-stuffing and CRC[3] error detection.

[1] Unshielded Twisted Pair, see Volume 1, Section 8.3.1.1.

[2] Non Return to Zero Inverted, see Volume 1, Section 8.3.6.4.

[3] Cyclical Redundancy Check, see Volume 1, Section 7.3.6.

3.1.2 FireWire[4]

[4] IEEE 1394

FireWire was invented by Apple Computers and Texas Instruments for local area networking, but has been standardised by the IEEE. It has a maximum data rate of 400 Mbit/s (or 800 Mbit/s for 1394b) and can support up to 63 devices on an external bus, which consists of up to 4.5 metres (without repeaters) of three pairs of STP cable. Devices typically have three Firewire ports and they can be connected in a cascaded star topology network. FireWire also distributes low voltage power to peripherals, which no longer require separate low voltage power supplies. FireWire supports plug and play attachment and all devices are hot pluggable.

It removes the possibility of contention by means of TDM[5] and allocates separate channels to each device connected to the bus. It is therefore ideally suited to carrying fixed bandwidth transmissions such as video. Sony have licensed FireWire for their camcorders and have branded it as iLink.

[5] Time Division Multiplexing, see Volume 1, Section 2.3.1.

FireWire uses Data-Strobe (DS) coding with NRZ[6] signalling. It transmits two signals on different pairs: a data signal and a strobe signal, which is produced by a bitwise XOR[7] operation on the data signal and an alternating clock signal. At the receiver the data and strobe signals are XORed to reconstruct the clock signal and hence maintain synchronisation.

[6] Non Return to Zero, see Volume 1, Section, 7.3.6.

[7] Exclusive OR (either one or the other but not both).

Unlike USB, it is a peer-to-peer protocol, although one station still has to act as a bus manager. This means that any two FireWire devices can communicate with each other, unlike USB where two slaves are unable communicate with each other directly.

Activity 3.1

Satisfy yourself that Data Strobe coding works by taking an arbitrary byte of data and produce a strobe signal by XORing the data with a byte containing an alternating pattern of 1s and 0s that represents a clock signal. Take the strobe signal and XOR it with the original data and check that you can regenerate the clock signal.

3.1.3 Infra-red Data Association (IrDA)

The Infra-red Data Association was founded in 1993 to develop and promote a standard for infra-red data communications. Its earliest members were Hewlett Packard, IBM and Sharp. It was designed for short-range digital communications between a PC and devices such as a keyboard, personal digital assistants, digital cameras, mobile phones and intelligent watches. The signal range is up to one metre, in normal light conditions and within a 30° cone from the transmitter. It can operate at speeds ranging from 9.6 kbit/s to 4 Mbit/s and uses three different modulation/coding techniques depending on the data rate. Below 115.2 kbit/s it uses asynchronous transmission (with start and stop bits), CRC-16 error checking and byte stuffing. It uses a coding scheme similar to RZ[8] with a zero represented by a pulse in the first part of the bit period and a one by no pulse. For data rates between 576 kbit/s and 1.152 Mbit/s, synchronous transmission with bit stuffing and CRC-16 error checking and a similar coding scheme is used. 4Mbit/s transmission is also synchronous but without the need of bit or byte stuffing. It uses CRC-32 error checking and a modulation technique known as pulse position modulation, which uses the position of pulses in time to represent the data.

[8] Return to Zero, see Volume 1, Section, 8.3.6.4

The IrDA standard include its own mandatory data link protocol (IrLAP), management protocol (IrMP) as well as optional transport protocols and other protocols to support particular functions.

3.1.4 Bluetooth

Bluetooth is a low-cost, low power short-range radio technology developed by the Bluetooth Special Interest Group made up initially of Ericsson, Intel, Nokia and Toshiba but now also containing Microsoft and many other members. It is has been designed to carry voice and/or data between mobile devices, such as mobile phones, hands-free headsets, digital cameras, lap-top computers and Personal Digital Assistants within a noisy radio environment in the unlicensed 2.4 GHz band. Bluetooth's signalling rate is 1 Mbit/s and it can support asymmetric data rates of 721 kbit/s and 57.6 kbit/s or symmetric rates of 432.6 kbit/s within a range of 10 metres. Bluetooth can also use this bandwidth to support up to three full duplex 64 kbit/s voice channels. Data is supported by Asynchronous Connection-Less (ACL) point-to-multipoint links and voice is supported by Synchronous Connection-Oriented (SCO) point-to-point links.

Up to eight active Bluetooth devices can form an ad-hoc network, called a piconet. One of the devices has to act as a master. Communication can only take place between a master and a slave. Two slaves cannot communicate with each other directly. A piconet can also support up to 255 non-active (or parked) devices, which can only become one of the eight active devices if there are less than eight active devices or if one of the active devices becomes parked. Up to ten piconets can exist and inter-operate in the same room as a scatternet formed when a slave of one piconet becomes a master of another and each bridging traffic between two piconets. Bluetooth uses TDMA[9],

[9] Time Division Multiple Access, see Volume 1, Section 7.3.8.

under the control of the master, to allocate channels to timeslots and FHSS[10] to transmit the timeslots and minimise the effects of interference. Each piconet follows a different frequency hopping pattern. Severe interference is likely to reduce the data rate, but will not stop Bluetooth from working altogether.

[10] *Frequency Hopping Spread Spectrum, see Section 3.3.5 of this Section.*

Bluetooth uses a 48-bit addressing scheme similar to that used by IEEE 802. It uses a FSK[11] modulation technique and CRC-16 error detection. For SCO, detected CRC errors result in frames being discarded. Three error correction techniques are used: ARQ[12] is used for ACL and two optional types of FEC[13] for both ACL and SCO. 1/3 FEC is applied to packet headers for both ACL and SCO and for some ACL data fields. It simply repeats bits three times. 2/3 FEC is used by some ACL and packet types and uses Hamming codes.

[11] *Frequency Shift Keying, see Volume 1, Section 8.3.6.3.*

[12] *Automatic Repeat Request, see Volume 1, Section 7.3.6.2.*

[13] *Forward Error Correction, see Volume 1, Section 7.3.6.2.*

Bluetooth has a Link Management Protocol to set up and control links and a Logical Link Control and Adaptation Protocol (L2CAP) which is its data link protocol and a protocol above this known as RFCOMM[14] which emulates a EIA-232 serial interface, so that software that normally runs over a serial interface can also run over Bluetooth.

[14] *Radio Frequency Communications*

Bluetooth incorporates a Service Discovery Protocol that allows applications to discover what functions are supported by other Bluetooth devices. Bluetooth creates a database of trusted devices following authentication which involves pairing the devices by entering PIN numbers on both.

It supports three modes of security: Mode 1 offers no security and is used by devices that have no secure applications. It can be used for transferring information such as business cards or calendar information between PDAs. Mode 2 offers service level security by means of authentication and encryption which can be varied for different applications. Mode 3 offers link mode security where all data links are encrypted from the time they are established using keys derived from the authentication process.

The physical and data link layers of Bluetooth is now being standardised as Wireless Personal Area Network (WPAN) by the IEEE 802.15 committee in conjunction with the Bluetooth Special Interest Group. It should be noted that Bluetooth uses the same frequency range as IEEE 802.11b[15] and that Bluetooth is likely to cause interference when used near WiFi hotspots. The IEEE have set up a task group to make recommendations that will help the two to co-exist.

[15] *Also known as WiFi, see Section 3.3.5 of this chapter.*

Activity 3.2

Find some information about the protocol layers used by Bluetooth and how they relate to the hybrid reference model.

Activity 3.3

Research the application functions that are supported by Bluetooth technologies on mobile phones.

3.2 Storage Area Networks (SANs)

SANs are similar in their range to local area networks but are dedicated networks that connect computers to data storage devices. Historically, computers have used external parallel buses such as the Small Computer Systems Interface (SCSI) to connect their data storage devices. With the advent of disk arrays and data warehousing, a requirement grew to build networks to allow many machines to access many storage devices. There are two distinct ways that this can be done. Network Attached Storage uses conventional LANs (or WANs) to carry a disk block protocol such as SCSI over the IP network layer protocol. This solution is good for a file server accessing whole files. But when the computer's operating system wants to transfer blocks of raw data efficiently, this is not an effective solution, in terms of latency and processing overhead. To do this a Storage Area Network running a specially designed protocol called Fibre Channel is preferred.

3.2.1 Fibre Channel[16]

[16] *ANSI X3T11*

Fibre Channel was invented by IBM and has been standardised by ANSI as an integrated set of standards. It was designed as a high speed protocol, supporting 100, 200, 400 and 1,200 Mbit/s over several kilometres of optical fibre (it can also be implemented over copper). It was specifically designed for use between computers and data storage devices to carry a number of block-based protocols, but is mainly used to carry SCSI. It can also carry IP.

Fibre Channel can support three different topologies. The simplest is a point-to-point topology between just two devices. The second is an arbitrated loop, which is a ring topology that supports up to 127 devices. Before transmitting a device must arbitrate to gain control of the loop. Contention is resolved by giving control to the device that has the lowest physical address. This device then has control of the loop for as long as required, but it is not allowed to arbitrate again until all the other devices have had the opportunity to arbitrate. The loop is often collapsed into a Fibre Channel hub, to simplify the wiring and protect the loop from breaking. The third topology is known as fabric, named after the cross-point switch it uses. This allows many devices (up to 2^{24}) to communicate with each other at the same time.

Three main classes of service are supported: Class 1 is a connection-oriented service equivalent to a dedicated fixed bandwidth physical connection; Class 2 is an acknowledged connectionless service and Class 3 is an unacknowledged connectionless service.

Fibre Channel supports flow control, which operates buffer-to-buffer or end-to-end using credit, which defines the rate at which devices can accept data. Each device must indicate its credit before establishing communications. Fibre Channel uses unique 64-bit port addresses known as WorldWide Names (WWNs), which are allocated to manufacturer's by the IEEE in a similar way to MAC addresses[17].

Fibre Channel uses 8B/10B[18] encoding and CRC-32 error checking.

[17] *Media Access Control addresses. See Volume 1, Section 7.3.2.*

[18] *8 Binary / 10 Binary, See Volume 1, Section 8.3.6.4.*

Activity 3.4

Search for some information about the protocol layers used by Fibre Channel and how they relate to the hybrid reference model.

3.3 Local Area Networks (LANs)

Local Area Networks became popular in the 1980s as a means to share data and expensive peripherals (such as printers) between PCs in offices. The IBM PC had just been launched (without any networking capability) and was being installed in large numbers by departments of many companies. Because the original PC was a stand-alone device, its introduction presented a number of problems. Firstly, data produced on one PC often had to be transferred to other PCs. The only way to do this was by means of a $5^1{}_4$ inch floppy disk. The actual transfer of data was not so much a problem, the real problem was ensuring that everyone using the data was working the most up-to-date version of it. There was therefore a need for central file storage, where everyone in the department could access the latest versions of data files. The other problem was that hard disk drives and peripherals such as printers were in those days rather expensive and a great deal of money could be saved if users were able to share them. This all led to the urgent need for a centralised file server and print server and a network over which files could be transferred. LANs were the solution to these problems they were soon being widely implemented by enterprises.

This itself created further problems as different departments within enterprises often implemented their own LANs and these were sometimes incompatible. A need arose to connect LANs within a building or across a campus. This could and sometimes had to be done using network layer protocols such as IP, but is often desirable to do this at the data link layer. Although it could be done by creating a single large LAN using bridges or swithes, it is often preferable to link LANs together using a higher speed LAN that is used solely for that purpose and does not have any stations directly connected to it. Such a LAN is called a Backbone LAN. A common design is to implement a separate LAN on each floor of a building using switches and then connect these switches to a Backbone LAN that runs in the vertical risers between the floors.

3.3.1 Ethernet

Ethernet was developed in 1976 by Xerox at their Palo Alto Research Center (PARC). The protocol was based on an earlier packet radio protocol (ALOHA) developed by the University of Hawaii and used to provide a packet switched radio data network between the Hawaiian Islands. It solved the multiple access problem using CSMA/CD[19] as its access method. Although Ethernet broadcasts over a wired medium (originally a coaxial cable) the contention problem between devices was the same as that experienced and solved by ALOHA. The developers of Ethernet therefore chose to implement a CSMA/CD protocol with a binary exponential back-off algorithm.

[19] Carrier Sense Multiple Access with Collision Detection, see Volume 1, Section 7.3.8.

Ethernet was marketed by a consortium of Digital, Intel and Xerox, known as the DIX consortium. It was the first protocol specifically designed for local area networking. Some earlier LANs had been built using HDLC[20]. Ethernet arrived on the market at just the right time for enterprises trying to share data and peripherals between PCs. Ethernet met this need and quickly became the de facto standard for local area networking. The mass production of Network Interface Cards (NICs) which could be installed in the expansion slots of PCs and servers led to a reduction in prices and the cost of installing LANs could easily be justified. As a result LANs, and Ethernet in particular, came into widespread use in offices. Along with other LANs, Ethernet uses 6-byte MAC addresses burnt into the NICs. The protocol supports unicast, multicast and broadcast addresses.

[20] High-level Data Link Control, see Volume 1, Section 7.5.

There are currently two variants of Ethernet frame formats[21] in use today. There are subtle differences between them regarding such things as the last byte of the preamble, but the main difference is in how the protocol indicates which network layer protocol is being carried.

[21] Ethernet II and IEEE 802.3, see Volume 1, Section 7.8

- **Ethernet II.** The original Xerox protocol (Ethernet I), which is now obsolete, had a length field but no field for de-multiplexing network layer protocols. The DIX consortium redefined the protocol as Ethernet II and replaced the length field with a type field which allows it to multiplex and de-multiplex the different network layer protocols that it can carry. The IETF[22] standardised on Ethernet II for carrying IPv4.
- **IEEE 802.3.** The Ethernet protocol standardised by the IEEE[23] is known as 802.3 after the sub-committee that specified it. IEEE 802.3 returned to using the length field and uses the SSAP[24] and DSAP[25] fields in the LLC[26] header for multiplexing and de-multiplexing different network layer protocols. A version of 802.3 without an 802.2 sub-layer known as

[22] Internet Engineering Task Force, see Volume 1, Section 3.3.

[23] Institute of Electrical and Electronic Engineers, see Volume 1, Section 3.3

[24] Source Service Access Point, see Volume 1, Section 7.7.

[25] Destination Service Access Point, see Volume 1, Section 7.7.

[26] Logical Link Control, see Volume 1, Section 7.7.

IEEE 802.3 Raw, exists and is used by Novell. A number of suppliers, such as IBM, Novell and Microsoft supported IEEE 802.3 but Ethernet II remained popular because of the IETF specifications. Because LLC only allows one byte for specifying the network layer protocol used, as opposed to two bytes supported by Ethernet II, a further sub-layer has been defined to support to allow IEEE 802.3 to support the full range of Ethernet II types. This is known as the Sub-Network Access Protocol (SNAP). It does this by using an LLC SSAP and DSAP codes AA16 which signifies that a SNAP header follows the LLC header.

Ethernet, suffers from having a number of different frame format standards, but all is not as bad as it first appears, as the different formats can coexist on the same LAN. This is because the lowest value of the Ethernet type field has been defined as 0600_{16} and the maximum length of an Ethernet frame is 1500 bytes which is coded as $05DC_{16}$. An Ethernet implementation is therefore able to determine whether the frame structure being used is Ethernet II or IEEE 802.3 by the value of the length/type field and hence the meaning of the field.

Ethernets can operate at different speed over different media which have different characteristics such as maximum cable lengths. Ethernet CSMA/CD specifications require that the round trip time on any segment is less than that required to transmit the minimum frame size of 512 bits to ensure that all collisions are detected. This will limit the maximum segment length for different data rates.

3.3.1.1 Ethernet designations

A naming convention (which has evolved over the years and is now very complex and confusing) is employed to describe these different types of Ethernet, as described in Table 3.2 below. You are not expected to learn the meaning of all the various parts of the designation, but you are expected to know the basic structure and the first four fields in the table which are commonly used.

Table 3.2: Ethernet naming convention

Field	Values	Meaning
Data Rate	1	1 Mbit/s
	10	10 Mbit/s
	100	100 Mbit/s
	1000	1 Gbit/s
	10G	10 Gbit/s
Transmission (after Data Rate)	Base	Baseband
	Broad	Broadband
	Pass	Passband (using bandpass filters, as in VDSL)
Max. Segment Length	2	185 m (approx 2 x 100 m)
(after Base or Broad)	5	500 m (5 x 100 m)
	36	3,600 m (36 x 100 m)
Media (after Base)	-C	Copper (Twinax cable)
	-F	Fibre
	-LH	Long Haul over dark fibre
	-T	Twisted pair (can also occur after Pass)
	-VG	Voice Grade twisted pair
Fibre Wavelength (after Base)	-E / -Z	Extra Long Wavelength(1550 nm)
	-L	Long Wavelength (1310 nm)
	-S	Short Wavelength (850 nm)

Table 3.2: Ethernet naming convention (cont.)

Field	Values	Meaning
Fibre Configuration (after Base)	-B	Bi-directional Point-to-Point
	-P	Passive Optical Network (Point-to-Multipoint)
Fibre Type (after -F)	B	Backbone
	L	Long
	P	Passive (without repeaters)
Twisted Pair Type (after -T)	S	Short Reach
	L	Long Reach
Coding (after -E, -L, or -S)	R	64B/66B
	W	WAN Interface Sublayer (64B/66B in SDH)
Coding (after -E, -L, -S, -B, -P, or -Z)	X	8B/10B
Fibre Mode (after -LR)	M	Multimode
Max. Segment Length (after BX, LX or PX)	10	10 km
Max. Segment Length (after PX)	20	20 km
Number of Fibres/Cable Pairs (after -T or X)	2 4	

Table 3.3 below gives examples of Ethernets (for reference) using the naming convention. You are not expected to know the details contained in this table, but you are expected to be able to recognise the speed and type of commonly implemented Ethernets (discussed in later this chapter) from their designations.

Table 3.3: Ethernet designations

Ethernet designation	Duplex	Media	Max segment length(m)	Comments
1Base5	Half	2 Cat 3 UTP	500	1 Mbit/s Star-based LAN using Twisted Pairs - Obsolete
2Base-TL	Full	1 Cat 3 UTP	2.700	2 Mbit/s Ethernet in the First Mile (EFM) over SHDSL
10Base5	Half	Thick Coax	500	10 Mbit/s Thicknet Coaxial Bus
10Base2	Half	Thin Coax	185	10 Mbit/s Thinnet Coaxial Bus
10Base-T	Full	2 Cat 5 UTP	100	10 Mbit/s Hubbed/Switched Ethernet
10Broad36	Half	Thick Coax	1,800	10 Mbit/s Coaxial Bus (3,600m between two stations) – Obsolete
10Base-FB	Half	2 MMFs	2,000	10 Mbit/s Backbone Ethernet using Multimode Fibre
10Base-FL	Full	2 MMFs	2,000	10 Mbit/s Ethernet using Multimode Fibre
10Base-FP	Half	2 MMFs	500	10 Mbit/s Ethernet using Multimode Passive Fibre
10Pass-T	Full	1 Cat 3 UTP	750	10 Mbit/s Ethernet in the First Mile (EFM) over a VDSL
100Base-BX10	Full	1 SMF	10,000	100 Mbit/s EFM over a single bidirectional point-to-point SMF
100Base-FX	Full	2 MMFs	2.000	100 Mbit/s Fast Ethernet over dual Multimode Fibre
100Base-LH	Full	2 SMFs	80,000	100 Mbit/s Fast Metro Ethernet – Long Haul over dark fibre
100Base-LX	Full	2 SMFs	5,000	100 Mbit/s Fast Ethernet using long wavelength LASERs
100Base-LX10	Full	2 SMFs	10,000	100 Mbit/s EFM using long wavelength LASERs
100Base-T2	Full	2 Cat 3 UTP	100	100 Mbit/s Fast Ethernet using 2 Twisted Pairs
100Base-T4	Half	4 Cat 3 UTP	100	100 Mbit/s Fast Ethernet using 4 Twisted Pairs
100Base-TX	Full	2 Cat 5 UTP	100	100 Mbit/s Fast Ethernet using Twisted Pair
100Base-VG	Full	4 Voice Grade UTP	100	Not strictly Ethernet - uses polling and not CSMA/CD
1000Base-BX10	Full	1 SMF	10,000	Gigabit EFM over a single bidirectional point-to-point SMF
1000Base-CX	Full	2 Twinax Cables	25	Gigabit Ethernet using Twinax cables - Obsolete
1000Base-LH	Full	2 SMFs	70,000	Gigabit Metro Ethernet – Long Haul over dark fibre

Table 3.3: Ethernet designations (cont.)

Ethernet designation	Duplex	Media	Max segment length (m)	Comments
1000Base-LX	Full	2 SMFs/MMFs	5,000	Gigabit Ethernet using Long Wavelength LASERs
1000Base-PX10	Full	1 SMF	10,000	Gigabit EFM over point-to-multipoint SMF using passive optics
1000Base-PX20	Full	1 SMF	20,000	Gigabit EFM over point-to-multipoint SMF using passive optics
1000Base-SX	Full	2 SMFs/MMFs	550	Gigabit Ethernet using short wavelength LASERs
1000Base-T	Full	4 Cat 5 UTP	100	Gigabit Ethernet using two Twisted Pairs
1000Base-ZX	Full	4 SMFs	70,000	Gigabit Metro Ethernet using extended wavelength LASERs
10GBase-CX4	Full	4 Twinax Cables	15	10 Gb Ethernet using Twinax cables - Obsolete
10GBase-ER	Full	2 SMFs	40,000	10 Gb Metro Ethernet over dark fibre
10GBase-EW	Full	1 STM-64 SDH cct	40,000	10 Gb Metro Ethernet over SDH WAN circuit (actually 9.95 Mbit/s)
10GBase-LR	Full	2 SMFs	10,000	10 Gb Metro Ethernet over dark fibre
10GBase-LRM	Full	2 MMFs	220	10 GbE over MMF using electronic dispersion compensation
10GBase-LX4	Full	2 SMFs	10,000	10 GbE (300 m over MMF or 10km over SMF using 4 WDM channels
10GBase-LW	Full	1 STM-64 SDH cct	10,000	10 Gb Metro Ethernet over SDH WAN circuit (actually 9.95 Mbit/s)
10GBase-SR	Full	2 MMFs	300	10 Gb Ethernet using short wavelength LASERs
10GBase-SW	Full	1 STM-64 SDH cct	300	10 Gb Metro Ethernet over SDH WAN circuit (actually 9.95 Mbit/s)
10GBase-T	Full	4 Cat 7 UTPs	100	10 Gb Ethernet using Unshielded Twisted Pair cables

3.3.1.2 Traditional Ethernets (10 Mbit/s)

Bus Ethernets

The original DIX Ethernet (10Base5[27]) used a thick (10mm) coaxial cable, similar to that used to connect TV sets to aerials, in a bus topology with stations being connected on a spur via a transceiver (sometimes called a vampire tap) that were tapped into the cable at intervals along its length. The physical interface used on the spur cable which could be up to 50 meters long was defined as the Attachment User Interface (AUI). As a result of the thick cable, it is sometimes referred to as Thicknet. More commonly though it is called 10Base5, with 10 indicating the data rate, Base indicating that it uses baseband signalling and 5 indicating that the maximum cable length is 500 meters.

[27] *IEEE 802.3*

An alternative and cheaper Ethernet is known as 10Base2[28] (sometimes called Thinnet or Cheapernet) cabling method used a thinner (5mm) coaxial cable and a 3-way connector of a type known as BNC[29] rather than a vampire tap. The PC could again be connected to the 3-way connector on the bus via a spur cable and a transceiver or more often an internal transceiver on the NIC card was used which allowed the 3-way connector to be directly connected to the NIC card at the back of the PC. The thin coaxial cable could only carry signals 185 meters and using the standard naming convention this type of Ethernet is known as 10Base2.

[28] *IEEE 802.3a*

[29] *Bayonet Neill Concelman (after its inventors) sometimes also erroneously called British Naval Connector.*

Both 10Base5 and 10Base2 were quite expensive to install. The coaxial cable was expensive and because it is so heavy and inflexible, it is difficult to run it through ducts.

Cabling for both 10Base5 and 10Base2 is expensive, bulky and quite difficult to install.

Traditional Ethernets use baseband signalling (the digital signals are not modulated onto another frequency) at 10 Mbit/s with Manchester encoding[30] at the physical layer.

[30] *See Volume 1, Section 8.3.6.4.*

Hubbed (or half duplex) Ethernets

Attempts were made to run local area networks using UTP cable in a star configuration. At the centre of the star was an active hub which when it receives a frame on port will broadcast it out on all the other ports (apart from the one from which it received the frame). A hub can therefore be regarded as a multi-port repeater. Thus the star network is able to emulate the standard bus network running the same Ethernet protocol and can be described as a physical star and logical bus network. This type of network running at 10Mbit/s is known as 10Base-T[31]. The star network is now the preferred topology for all types of Ethernet.

[31] IEEE 802.3i

Switched (or full duplex) Ethernets

A further improvement can be made by replacing the hub by a switch. The switch is an intelligent device that examines a frame's destination MAC address and only forwards the frame out of the one port to which the device with that physical address is attached. It learns the physical addresses of each device from the frames that it receives in the same way as a bridge. If it receives a frame with a destination address that it does not recognise, then it broadcasts the frame out of all its ports, apart from the one on which it was received. Using a switch rather than a hub greatly improves the efficiency of an Ethernet, as it converts a single collision domain, where transmissions from any device can collide with transmissions from any other device, to multiple collision domains, where there is usually only one device able to transmit on the physical medium. As a result switched Ethernets can support much higher throughput than hub-based or bus Ethernets.

Hubbed and switched Ethernets often provided dedicated media for each direction of transmission. This means that for switched Ethernets no collisions can ever occur and there is no requirement for CSMA/CD. Attached devices can receive and transmit at the same time and hence can operate in full duplex mode. Also the segment length limit requiring the RTT[32] to be less than the time taken to transmit 512 bits is no longer required and hence segment lengths can be longer. If the same media is shared for both transmitting and receiving (as is the case with a hub) then a single device will be unable to receive and transmit at the same time and will hence operate in half duplex mode. Full duplex Ethernets also double the effective capacity as 10 Mbit/s can be carried simultaneously in both directions giving an overall maximum theoretical capacity of 20 Mbit/s.

[32] Round Trip Time.

Switches can operate in different modes.

- Store and Forward mode is where the whole frame is stored in a buffer and the CRC field is checked before forwarding.
- Cut-through mode is where the switch starts to forward the frame as soon as the destination physical address has been received.
- Fragment Free mode is where the switch starts to forward the frame after the first 64 bytes have been received. Fragments of 64 bytes or less (also called runts) are often created by collisions and by checking that no collision has occurred while the first 64 bytes of a frame have been received the switch can ensure that no fragments are generated.
- Adaptive mode is where the switch continually monitors the number of collisions and errors and determines which mode is best to use.

Clearly, cut-through mode is superior in its delay performance, but only if the number of collisions and hence errors are low. Cut through mode is ideal with full duplex Ethernets where collisions never occur. Store and Forward mode will delay every frame, but will ensure that frames containing errors are not forwarded and will perform best in half-duplex Ethernets that have a high collision rate. Fragment Free is a compromise between the two, while Adaptive mode can dynamically adjust the mode depending on the collision rate.

3.3.1.3 Fast Ethernet (100 Mbit/s)

Fast Ethernet works at 100 Mbit/s and was originally intended for use in backbone LANs, but is now used routinely in LANs supporting servers or many PCs requiring high bandwidth services. The upgrade path from 10Base-T to 100Base-TX[33] or 100Base-T2[34] or from 10Base-FL[35] to 100Base-FX8[36] is relatively straightforward as the same cable can be used, as they support the same maximum segment lengths. 100Base-TX requires an extra cable pair, but Cat 3 cable always comes with two pairs, only one of which is used in 10Base-T. Only the Network Interface Cards and switches need to be changed or reconfigured to work at a higher speed.

[33] IEEE 802.3u
[34] IEEE 802.3y
[35] IEEE802.3u
[36] IEEE802.3u

Fast Ethernet supports the same frame structure, minimum and maximum frame lengths and addressing as traditional Ethernet as well as the MAC layer including the CSMA/CD access method (if appropriate). This means that both types of Ethernet are fully compatible and that a switch can actually support both speeds. Typically lower specification PCs can be connected to the switch at 10 Mbit/s and servers and high specification PCs and trunk links to other switches can be connected at 100 Mbit/s. The speed mismatch will not usually cause any problems because frames received at high speed ports are likely to be destined for many low speed ports and it is unlikely that any 10Mbit/s port will ever need to receive frames at a higher data rate than that supported. On most switches, it is possible to configure switch ports for autonegotiation where the speed at which a device is operating is detected by the switch and configured accordingly. This allows a single device to be configured to work at either speed without having to reconfigure the switch.

100Base-T Ethernets do not use Manchester encoding but uses various alternative more modern coding methods, such as MLT-3[37] for 100Base-TX, PAM5[38] for 100Base-T2, 8B/6T NRZ[39] for 100Base-T4 and 4B/5B[40] NRZI for 100Base-FX.

100Base-VG[42] was an attempt to produce a single LAN standard for 100Mbit/s backbone LANs. It used four pairs of Cat 3 UTP (also known as voice grade) and hence the initials. It is also sometimes called 100VG-AnyLAN as it was designed to carry both Ethernet and Token Rings frame types on the same hub. It was originally proposed by HP[43] and the IEEE started to standardise it as an 802.3 protocol, but because it used a polling access method to overcome segment length limitations rather than CSMA/CD, it was decided to progress the standard via a new committee (802.12). It was not successful and is now practically extinct as CSMA/CD protocols were developed that could support 100m segment lengths.

[37] Multi-Line Transmission – 3-level, see Volume 1, Section 8.3.6.4.
[38] Pulse Amplitude Modulation – 5 state.
[39] Non-Return to Zero, see Volume 1, Section 8.3.6.4.
[40] 4 Binary / 5 Binary, see Volume 1, Section 8.3.6.4.
[41] Non-Return to Zero Inverted, see Volume 1, Section 8.3.6.4.
[42] IEEE 802.12
[43] Hewlett Packard

3.3.1.4 Gigabit Ethernet (1 Gbit/s)

Gigabit Ethernet (GbE) supports speeds of 1 Gbit/s and was designed for use over backbone LANs. It is required because many organisations upgraded their LANs to Fast Ethernet which put huge traffic loads on their backbone LANs. GbE is also required for server farms and may ultimately be required at the desktop for applications that require very high bandwidth. GbE supports the same frame structure, maximum frame lengths and addressing as traditional and Fast Ethernet as well as the MAC layer including the CSMA/CD access method (if appropriate), but the minimum frame size has been altered. The rule that restricts segment lengths so that the RTT is no more than the time taken to transmit 512 bits will result in the maximum segment length being impractically short. The 512 bit limit is therefore changed to 4096 bits and frames are padded out so that they are always at least 4096 bits (512 bytes) long. This is known as carrier extension. Another change introduced with GbE allows a station which has a number of short frames awaiting transmission to burst them consecutively as a single transmission, to avoid the overhead of carrier extension.

Although the GbE standard has gone to a lot of trouble to implement CSMA/CD for half duplex operation to be backwards compatible, it is rarely implemented and most GbE is full duplex, so the limits on maximum segment length are governed by signal strength rather than collision detection issues.

GbE can be implemented over copper or fibre. Early copper implementation, such as 1000Base-CX[44] used a special twin coaxial cable (known as twinax) and 8B/10B[45] coding, but modern implementations use 1000Base-T[46] which uses 4 Cat 5 UTPs and PAM5 coding.

Fibre implementations of GbE cannot use LEDs[47] as they cannot be switched on and off quickly enough, so LASERs[48] have to be used. These can use short (1000Base-SX[11]), or long (1000Base-LX[11]) wavelengths and also use 8B/10B coding.

44 IEEE 802.3z

45 8 Binary / 10 Binary, see Volume 1, Section 8.3.6.4.

46 IEEE 802.3ab

47 Light Emitting Diode

48 Light Amplification by Stimulated Emission of Radiation

3.3.1.5 10 Gigabit Ethernets (10 Gbit/s)

The main driver for 10 Gigabit Ethernets (10GbE) is the need to switch large volumes of Internet traffic between ISPs across a shared LAN in a building that acts as an interconnection point between ISPs[49]. 10GbE will also be useful for large corporations who have a requirement to switch large amounts of data between servers and their Intranets or the Internet.

49 Such as the London InterNet eXchange (LINX)

Unlike other Ethernet standards 10GbE has from the outset been designed to support MAN and WAN communications. Also, unlike other Ethernet standards, it always runs in full-duplex point-to-point mode, so there is no possibility of any collisions and hence no need for collision detection. With 10GbE, Ethernet has finally freed itself from its CSMA/CD access method origins.

10GbE uses short (S) or long (L) wavelength LASERs with 8B/10B (X) or 64B/66B coding (R). 10GBase-T uses 4 Cat 7 UTPs and PAM-12 line code.

3.3.1.6 Metro Ethernet

An important requirement for many organisations is to interconnect their LANs over a metropolitan area network. This was often done in the past by routing over private circuits using HDLC, or using public networks such as frame relay, or ATM[50] or SMDS[51]. Only the latter two could support LAN speeds and network operator charges for these services were high. Using routers and a different data link protocol also introduced inefficiencies and reduced throughput. An ideal solution would be to buy fibre capacity from the network operator and run Ethernet protocols over the fibre between two LAN switches. Network Operators offer high speed LAN extension services but their tariffs normally increase significantly with speed, but some operators are willing to sell 'dark fibre' where optical connectivity is provided between two sites without LASERs and detectors. Customers can then buy their own LASERs and detectors to 'light' the fibre at whatever data rate they require. This allows customers to set up their own high speed LAN interconnections where the costs do not increase so severely as data rates increase.

50 Asynchronous Transfer Mode, see Section 6.5 of this volume.

51 Switched Multi-megabit Data Service, see Section 4.1.1 of this volume.

100Base-LH can provide long haul connectivity up to 80 km while 1000Base-LH and 1000Base-ZX can reach up to 70 km. 10GBase-LR (using 64B/66B coding), 10GBase-LX4 (using 8B/10B coding) and 10GBase-LW can reach 10 km while 10GBase-ER (using 64B/66B coding) and 10GBase-EW can reach 40 km. 10GBase-LW and 10GBase-EW both operate over WANs using 64B/66B coding and encapsulate Ethernet frames within SDH[52] frames. They actually operate at about 9.95 Mbit/s to match the SDH STM-64 speed.

[52] Synchronous Digital Hierarchy, see Section 4.2.1.1 of this volume.

3.3.1.7 Ethernet in the First Mile (EFM)

EFM is a recent development of the protocol to run over broadband copper or fibre in the first mile (or the local loop) between customer premises and the public network. It makes sense to implement Ethernet here as Ethernet provides a simple and familiar interface which can operate at a range of speeds and for which standard Network Interface Cards exist and because of mass production they are very cheap. EFM is not only required for residential access, but is also required for business purposes to replace existing private circuit solutions and to provide integrated high speed access to various network services.

A number of solutions have so far been proposed for various media and topologies. For point-to-point services over copper at lower speeds 2Base-TL (for Twisted Pair Long Reach, but also sometimes called 2Base-S for SHDSL[53]) provide the equivalent of a 2Mbit/s private circuit to carry data over SHDSL circuits of up to 2.7 km. Currently a small business may lease a number of Nx64 kbit/s circuits to provide data communications to other sites or access to public data networks. These could all be carried over a single 2Base-TL circuit at greatly reduced cost.

[53] Symmetric High-bit-rate Digital Subscriber Line.

For higher speed point-to-point services at 10 Mbit/s over copper, 10Pass-T (also sometimes known as 10Pass-V for VDSL[54]) can support 10 Mbit/s over 750m. VDSL, like ADSL[55] and unlike SHDSL, does not use baseband signalling as it splits off voice signals by means of a low pass filter, so it is described as using passband signalling. The maximum segment length is insufficient to reach a network operator's premises in most cases, so the operator has to install an Optical Network Unit in a street cabinet within reach of the customer's premises.

[54] Very-high-speed Digital Subscriber Line

[55] Asymmetric Digital Subscriber Line.

For 100 Mbit/s point-to-point services over fibre 100Base-LX10 can reach up to 5 km and for 1 Gbit/s point-to-point services over fibre 1000Base-BX10 can reach 10 km.

Finally, technology is currently being developed to offer point-to-multipoint services over a Passive Optical Network (PON). With this technology multiple sites are connected to a single tree topology fibre network and sites can communicate with each other using a particular wavelengths of light. The switching in the network is therefore carried out by users selecting different wavelengths and the network itself is completely passive. Because the network is point-to-multipoint it is ideal for broadcasting and multicasting. This capability is also useful for data communications and Ethernet protocols have been developed for use over Passive Optical Networks. 1000Base-PX10 and 1000Base-PX20 can carry Gigabit Ethernet over 10 and 20 kms respectively.

Activity 3.5

Find out the designations (E.g. 10Base-T etc.) of the all the different types of Ethernet are used by your institution or place of work to support the desktop and in the backbone. Find some information about all the Ethernet types you discover.

3.3.1.8 Virtual LANs (VLANs)

LANs were historically set-up and administered on a departmental basis, but companies tend to reorganise and move employees around quite regularly, which causes a problem for the network managers who sometimes have reconfigure ports on switches and if the switches are not collocated then some new cabling will be necessary. Another issue is security. Some departments, notable Human Resources, Finance and Marketing departments hold sensitive information and are concerned that this information can only be accessed by members of their own departments. There has therefore been some resistance to merging or even bridging departmental LANs. Another problem results when LANs are merged or bridged is that of **broadcast storms**. A broadcast storm occurs when there is a sufficiently large number of broadcast messages on a LAN which collide with other messages and can make the LAN virtually unusable. Sometimes a faulty NIC card can cause a broadcast storm, but often they are caused by protocols that use an excessive amount of broadcasting. Clearly the larger the LAN (including a multi-switch LAN), the bigger this problem can be.

A solution to all of the above problems is to use Virtual LANs (VLANs). A VLAN can be thought of as a LAN that is configured by software on a LAN switch. The switch is configured to support a number of workgroups. Each workgroup is assigned its own VLAN. If an employee changes workgroups or moves to a different part of the building then any reconfiguration required can be done in the VLAN configuration rather than by reconfiguring ports on the switches.

Traffic between VLAN cannot be switched, but it can be routed. Routers however provide security functions that allow very tight control on what traffic can be routed which is not possible with switches. Network managers can therefore easily control traffic going to particular VLANs, by only allowing approved access and, by default, only members of the appropriate workgroup will be able to access systems on the VLAN. Furthermore, ensuring that all traffic between VLANs is routed and not switched, means that broadcast storms will not be forwarded from one VLAN to another, thus the larger LAN can be protected from broadcast storms on individual VLANs.

The network manager can define VLAN workgroups by:

- defining which switch ports belong to which VLAN (Port-based VLAN)
- defining which MAC addresses belong to which VLAN (MAC-based or Layer 2 VLAN)
- defining which IP addresses or which IP multicast address group belong to which VLAN (IP-based or Layer 3 VLAN)
- defining which multicast IP address belongs to which VLAN
- defining which applications can use the VLAN by means of the transport layer port number.

In a multi-switch LAN, when a switch receives a frame over a trunk from another switch, it must be able to tell which VLAN the frame relates to. There are three ways in which this can be done:

- Each switch can maintain a table of their VLAN members and can broadcast this table to the other switches periodically.
- TDM can be used on the trunk so that, for instance, frames for VLAN 1 always appears in timeslot 1.
- Each frame can be tagged with an additional header field that identifies the VLAN. This requires a change to the data link layer protocol or an additional sub-layer, but is the most common way of trunking VLANs.

Activity 3.6

Check whether Virtual LANs are used in your institution or place of work and if so, find out the names/identities of all the VLAN workgroups.

3.3.2 Token Ring[56]

[56] IEEE 802.5

The Token Ring network was invented by IBM for use with its office automation products and the specifications were handed over to the IEEE for standardisation.

The Token Ring as its name implies uses a ring technology with a token passing access method. It can operate at 1, 4 or 16 Mbit/s and uses Differential Manchester[57] encoding and the same 6-byte addressing scheme as Ethernet. The ring is connected using segments of Cat 3 UTP between stations. Each station, when powered, acts as a repeater for the signals circulating around the ring. If a station is switched off then a relay on the Network Interface Card operates to make a passive electrical connection through the card. This ensures that signals can pass through a station that is switched off. One of the stations has to act as an Active Monitor which must continually check that the ring is operating correctly and particularly that the same frame is not circulating the ring more than once and that the token is being circulated. If the Active Monitor detects that a frame has circulated more than once, it removes it from the ring. If it does not see a token before a timeout expires, it assumes that the token has been lost and creates a new one.

[57] *See Volume 1, Section 8.3.6.4.*

The token is a special 3-byte bit pattern that is passed from station to station around the ring. When a station wishes to transmit, it has to wait for the token to be received. It then takes the token off the ring, transmits its frame and the station which has the physical address indicated in the destination address field of the frame header retrieves the frame, performs a CRC check on it and if successful, it sets some bits in the frame status field to indicate that the frame has been successfully received by the correct station. When the frame returns to the transmitter with these bits set, it is removed from the ring and the token is then forwarded to the next station on the ring.

The token passing access method used ensures that a station is able to transmit a frame within a fixed time and that each station is treated fairly, apart from the use of priority and a scheme called early token release that is only used on 16 Mbit/s rings to improve efficiency. An upper bound on access time and fairness is not possible with Ethernet protocols, as under severe congestion the CSMA/CD algorithm means that an arbitrary time elapses before a frame is successfully transmitted without a collision. In the worst case, the frame may be discarded after the maximum number of attempts has been exceeded. Token Ring is therefore a deterministic protocol while Ethernet is non-deterministic.

Token Ring networks are vulnerable to cable breaks which will cause catastrophic failures. The danger of a cable break can be reduced by changing to a star topology (similar to that used by twisted pair Ethernet). The ring is effectively collapsed into a central hub known as a Multi-Station Access Unit (MAU) with spurs of twisted pair cables going out to each station and back to the hub. The relay that detects when a device is switched off is located in the hub and it will also be able to detect a cable failure and will close to bypass the spur going out to the device.

In many ways Token Ring is technically superior to Ethernet, but it has not been as successful in the marketplace. This is largely because Ethernet had a head start over Token Ring and is slightly simpler and cheaper to implement. As in most things success breeds success and the large number of Ethernets being implemented led to economies of scale and hence cheaper equipment. There may also have been some suspicion regarding IBM's involvement in its development, as at the time companies were attempting to escape the stranglehold of IBM's SNA[58] networking protocols.

3.3.3 Fibre Distributed Data Interface (FDDI)[59]

Unlike virtually all other popular LAN protocols, FDDI was not standardised by the IEEE. It was in fact standardised by ANSI[60] and ratified by ISO[61]. It was designed as for use in LAN backbones. It is highly compatible with other LAN protocols, such as Ethernet and Token Ring when connected via bridges. It operates at 100 Mbit/s uses a dual fibre ring topology with data circulating in different directions around each ring. It can use mono-mode or multi-mode fibre. Within a building multi-mode is likely to be used for cost reasons, but mono-mode may be used over a campus as it supports longer cable lengths. FDDI uses a token passing access method, similar to token ring, 4B/5B encoding, NRZI signalling and IEEE 802 MAC addressing.

At any one time, one ring is deemed to be the primary ring and the other the secondary ring which is always on hot-standby in case the primary ring breaks. If both rings break in the same place as in Figure 3.1 below, then the FDDI nodes either side of the break can loop the fibres together to make a single double size ring and thus maintain communications.

Figure 3.1: FDDI Recovery from a single cable break affecting both fibres

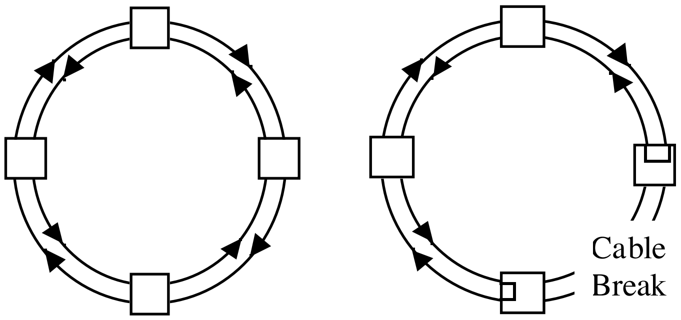

As with Fibre Channel and Ethernet hubs and Token Ring MAUs, the FDDI ring is best collapsed into a box known as a concentrator with fibre spurs going out to individual stations and back that can be bypassed when the station is switched off or when there is a cable fault. The need for this is even greater for FDDI as optical power limitations mean that it would not be feasible to bypass a station on a large ring, but it can be done when the ring is collapsed into a concentrator.

A copper version of FDDI, known as CDDI, has also been defined that operates at 100 Mbit/s over a dual collapsed ring with spurs of up to 100m of STP[62] going to and from the desktop. This was specified to allow FDDI to be extended to the desktop using cheaper technology than FDDI.

Neither FDDI and CDDI have enjoyed long-term success in the market place, although they are resilient and deterministic in the maximum time they take to transmit a frame. The protocols are complex and the technology is

[58] *Systems Network Architecture*

[59] *ANSI X3T12*

[60] *American National Standards Institute, see Volume 1, Section 3.3.*

[61] *International Organization for Standarization see Volume 1, Section 3.3.*

[62] *Shielded Twisted Pair see Volume 1, Section 8.3.1.1.*

expensive and could not compete with Ethernet which could not only be developed to offer equivalent speeds at less cost, but could also eventually provide an upgrade path beyond 100 Mbit/s.

3.3.4 Wireless LANs (WLANs)

WLANs have become increasingly popular in the last few years. The main reasons for this are that it is very expensive to cable a building and companies often move their offices to new buildings. Every time they do so, there is a large cost. Secondly a large number of employees now own lap-top computers and expect to be able to use them anywhere within the building. There is a considerable extra cost of providing spare LAN ports for lap-tops and a certain amount of hassle for users in finding these ports and connecting their lap-tops to them. Even then, the user cannot roam around the building and maintain a continual connection with the network. Users have experienced the flexibility that has come with mobile phones and expect something similar for data. Wireless LANs have been standardised by the IEEE 802.11 Committee who have defined two service sets.

- The Basic Service Set (BSS) consist of a number of wireless stations and a single Access Point (AP) or base station which has a connection to the fixed network, usually in the form of a wired IEEE LAN.
- The Extended Service Set (ESS) also includes multiple BSSs with their Access Points and stations. The APs are connected by a Distribution System (DS) which can be any IEEE LAN (wired or wireless).

There are three types of station supported:

- No Transition Mobility, where the station always remains in the same BSS
- BSS-Transition Mobility, where the station can move from one BSS to another but must stay within the same ESS
- ESS-Transition Mobility, where the station can move between ESSs.

IEEE Wireless LANs are based on Ethernet, but they cannot use the CSMA/CD access method. This is because listening for collisions while transmitting is not feasible. It would require duplicate radio circuitry which would double the cost of the equipment. Wireless stations therefore have to switch between transmitting and receiving and thus work in a half duplex mode. Also, even if the station could transmit and receive at the same time, it is quite possible that it would not be able to detect collisions caused by other stations in the BSS due to obstacles that cause fading and to the attenuation of the signals. For the above reasons another access method must be devised and the method chosen is CSMA/CA where CA stands for Collision Avoidance. With this method, a station wanting to transmit follow the following procedure:

- Before transmitting, it puts its radio into receive mode and senses the power levels at the carrier frequency.
- If it does not sense a transmission, it waits for a short period and transmits a Request to Send (RTS) frame. The RTS frame includes an indication of the length of time that will be required to transmit the frame.
- If it detects another station is transmitting it will carry out a binary exponential back-off and sense the carrier frequency again later, until it detects no other transmission, after which it will send its RTS frame.
- The receiver will respond to the RTS frame with a Clear to Send (CTS) frame which also includes the indication of the length of time it will take to transmit the frame.

- If no CTS is received, a collision is assumed and the transmitter backs off.
- Once the transmitter has received the CTS, it knows that it has access to the channel and can transmit its data frame after a short wait.
- The collision avoidance method works because all other stations in the BSS will see the RTS or the CTS and know how long the channel will be occupied for. They can therefore delay their next attempts to transmit accordingly.
- The transmitter must then await an acknowledgement from the receiver that the frame has been received. The acknowledgement is necessary, as the transmitter would have been unable to detect any collision while it was transmitting the frame.

The CSMA/CA access method is mandatory for all 802.11 WLANs. It is known as the Distributed Coordination Function (DCF). It is ideal for carrying data traffic and scales well as the number of stations on the WLAN grow. Like Ethernet it is non-deterministic and there is no upper limit to the time that it can take to transmit a frame, or a guarantee that a frame will be transmitted at all. For certain applications, such as voice, it is desirable to have an upper bound on frame transmission time. 802.11 allows this to be done by an optional polling method called Point Coordination Function (PCF). With PCF, the Access Point eliminates contention by polling each station in turn and therefore an upper bound to frame transmission time can be guarantied. PCF and DCF can coexist on the same WLAN by allocating different timeslots to each method. PCF does not scale as well as DCF, as the single Access Point in a BSS cannot poll an ever increasing number of stations.

802.11 defines a number of services offered by the Distribution System. These are:

- association, which allows the station to associate itself with an Access Point
- disassociation, which allows the station to remove itself from an association with an Access Point
- re-association, which allows a station to change its BSS by associating with a different Access Point which can be done without any loss of data during the handover
- distribution, which forwards frames to the correct destination via the Distribution System
- integration, which converts frame formats, where necessary, between the 802.11 protocol and the protocol of the Distribution System.

Other services are offered by the stations. These are:

- authentication, which allows stations to establish their identities by responding to challenges using their secret keys prior to transmitting data to each other
- de-authentication, which allows an authentication relationship between two station be terminated
- privacy, which allows stations to encrypt and decrypt the data they are transmitting using the RC4 algorithm
- data delivery, which allows data to be transmitted unreliably between stations.

The original 802.11 specification supported relatively low data rates (1 or 2 Mbit/s) using three alternative transmission schemes. These were:

- Direct Sequence Spread Spectrum (DSSS) operating in the 2.4 GHz band, where each bit from each transmitter is replaced by a sequence of bits (known as the chip code) which has the effect of spreading the signal across a large frequency range and increasing its immunity to noise.
- Frequency Hopping Spread Spectrum (FHSS) also operating in the 2.4 GHz band, where the carrier frequency hops over a range of frequencies according to an agreed sequence.
- Infrared using infra-red light which cannot pass through solid objects and is affected by the presences of sun-light. It was therefore not a popular choice of technology.

The original specifications have now been supplanted by separate specifications that offer higher data rate. The main ones are:

- IEEE 802.11a, which offers between 6 and 54 Mbit/s in the 5 GHz band and uses Orthogonal Frequency Division Multiplexing (OFDM), which is similar to FDM except all the sub-bands are used by the same transmitter at the same time
- IEEE 802.11b, which offers 1, 2 5.5 and 11 Mbit/s in the 2.4 GHz band and uses High Rate Direct Sequence Spread Spectrum (HR-DSSS) which is similar to DSSS but operates at lower frequencies and can hence support a higher data rate
- IEEE 802.11g, which offers 54 Mbit/s in the 2.4 GHz band and also uses OFDM.

Activity 3.7

Find some information about the protocol layers used by IEEE 802.11 LANs and how they relate to the hybrid reference model.

Specimen examination question

(a) State whether each of the following statements is true or false and, if false, correct the statement:

  i. A Bluetooth piconet consists of a master and up to 7 active slaves, but the slaves cannot communicate with each other directly.

  ii. IN IEEE 802.11 LANs, authentication is a service provided by the Distribution System which allows the station to authenticate itself to an Access Point.

  iii. Token Ring operates at 1, 4 or 16 Mbit/s and uses Manchester encoding.

  iv. Implementing VLANs not only provide better security but they reduce the effect of broadcast storms.

(b) Describe the main differences between USB and FireWire.

(c) Describe the four switching modes of an Ethernet switch and indicate under which conditions each is best used.

(d) Explain the advantages of implementing VLANs rather than using a single large multi-switch LAN.

(e) Describe how the CSMA/CA protocol works so that a frame can be transmitted without a collision.

Learning outcomes

At the end of this chapter, you should be able to:

- identify the approximate ranges over which PANs, SANs, LANs, MANs and WANs operate
- outline how the USB and FireWire protocols operate

- outline how the IrDA and Bluetooth protocol works
- outline how the Fibre Channel protocol works
- distinguish between and describe the various main types of Ethernet from their designations
- describe how Ethernet has evolved to higher speeds and to wider areas
- describe why VLANs are required and how they work
- outline how the Token Ring and FDDI protocols work
- describe the drivers for and main features of Wireless LANs and the CSMA/CA access method.

Appendix B: Model answers and hints

Chapter 2

(a) True or False

i. FALSE – a problem child has high market growth.

ii. TRUE.

iii. FALSE – it has done so in many countries by means of subsidiaries.

iv. TRUE.

(b) See bullet point list in Section 2.1.

(c) The main difference relates to the lack of an incumbent in mobile markets. This means that the markets are more dynamic, innovative and competitive. Because of the lack of competition, this fixed market is more heavily regulated than the mobile market. The fixed market is also much nearer being fully saturated, so high growth is harder to achieve than it is in the mobile market. The mobile market can appeal to a younger customer base than the fixed market. It is generally only householders who buy fixed network services, but mobile services can be bought by children, young people and students.

(d) See Section 2.2.4, first paragraph.

(e) See Section 2.2.5.

Chapter 3

(a) True or False

i. TRUE.

ii. FALSE – authentication is a service provided by the station.

iii. FALSE – it uses Differential Manchester encoding.

iv. TRUE.

(b) See Sections 3.1.1 and 3.1.2.

(c) See bullet point list in Section 3.3.1.2.

(d) See Section 3.3.2, first paragraph.

(e) See Section 3.3.5, third bullet point list.