Computer Multimedia and Networking Hardware System

John Martin

COMPUTER MULTIMEDIA AND NETWORKING HARDWARE SYSTEM

COMPUTER MULTIMEDIA AND NETWORKING HARDWARE SYSTEM

John Martin

BIBLIOTEX
Digital Library

Computer Multimedia and Networking Hardware System
by John Martin

Copyright© 2022 BIBLIOTEX

www.bibliotex.com

All rights reserved. No part of this book may be reproduced or used in any manner without the prior written permission of the copyright owner, except for the use brief quotations in a book review.

To request permissions, contact the publisher at info@bibliotex.com

Ebook ISBN: 9781984664112

Published by:
Bibliotex
Canada
Website: www.bibliotex.com

Contents

1

Applications in Networked Multimedia

Multimedia can roughly be defined as a technology that enables humans to use computers capable of processing textual data, audio and video, still pictures, and animation. Applications range over entertainment, education, information provision, design *e.g.* CAD/CAM, co-operative working such as video conferencing, application sharing, remote working and virtual reality experiences. Multimedia applications for computers have been developed for single computing platforms such as the PC, Apple Mac and games machines. The importance of communications or networking for multimedia lies in the new applications that will be generated by adding networking capabilities to multimedia computers, and hopefully gains in efficiency and cost of ownership and use when multimedia resources are part of distributed computing systems. Widening of access to multimedia sources and potential markets in multimedia, video and information are commercial driving force for networking multimedia.

The reality of networking multimedia is that :-

The characteristics of multimedia make heavy demands on storage and transmission systems.

Data compression can be used to reduce the demands of multimedia, particularly of video and audio on these systems, but usually at the expense of some loss in the detail compared with the source and at extra cost.

The ways in which users or participants in multimedia sessions access multimedia or connect with others have important consequences for the storage and transmission systems. For instance multimedia learning material can be accessed directly from a server during a class or downloaded to student machines prior to a session. The demands on a connecting network are very different in each access mode.

The cost of transmitting multimedia information will determine the pace of development of networked multimedia applications.

The availability of standards for multimedia networking, particularly for inter-working between applications, the development of networked applications, and interworking between networks are essential to reduce the complexity and level of skill required in using multimedia.

An Example

Using your desktop conferencing from Fujitsu which works on ISDN at 64 Kbps, and the Olivetti PCC video conferencing, you want to do some work with a colleague at Imperial College in London who fortuitously has a Super JANET Asynchronous Transfer Mode (ATM) connection at 34 Mbps, and has a workstation with UNIX based desktop conferencing and a codec for H.320 compatible video conferencing. Hard luck, you can't. While the video conference systems are compatible, both use H.320, the ISDN network cannot connect right through the ATM network. Also the desktop conferencing systems use different standards for sharing applications, whiteboards etc. and would not inter-work even if directly connected.

USER REQUIREMENTS FOR MULTIMEDIA

HUMAN-COMPUTER INTERFACE

The standards of reproduction for computers which are desirable have been set by the publishers of books, music, Walt Disney cartoons and television producers. With the development of High Definition TV and beyond, it is likely that there will be a continual increase in the demands placed on computer based multimedia systems.

The current PAL standard in the UK delivers video in 625 lines at 25 frames/sec. High Definition TV delivers video in 1250 lines with a higher horizontal resolution at 25 frames/

sec and requires about five times the information rate as the current PAL system.

Multimedia applications like any other application, appliance or tool, benefit from being easy to use, with minimal training or self learning. The need for a well designed human-computer interface, which may be screen or audio based is well accepted.

Access, Delivery, Scheduling and Recording

Television channels can be changed at the touch of a button. On demand access times to computer information need to be below one second to be usable in real time. Alternatively the delivery of information at a later time is acceptable if it can be scheduled, as in a TV broadcast schedule, or a first class postal letter. Scheduling the delivery of multimedia information has not been widely implemented. Scheduling can have advantages for users over on demand delivery. In a learning situation times can be defined for class attendance by a lecturer. In open learning situations learners can control their programme by requesting a multimedia unit at a convenient time. Just as we can record a TV film on a VHS recorder, some multimedia computer users will wish to record a film, session, or learning experience for future reference.

Interactivity

Interactivity, meaning the ability to participate in a video or audio process on a computer, by changing its behaviour or appending comments has become very important in multimedia. Some of this popularity stems from the perception of computer games as being enjoyable because they are interactive, and some from work done in education which shows that some types of learning becomes easier, and is retained more permanently if the learner participates in some way with the learning material.

Computer based multimedia needs the same degree of interactivity that a school exercise book, or a laboratory experiment has in order to remain credible as a learning medium. The generation of computer based virtual reality is an extension of this process. The incorporation of interactivity is really the job of the application designer. The incorporation of interactivity is assisted if the network is capable of two way communication, and for some applications the sense of interactivity is aided by the ability to deliver a moving picture, or a sound very quickly, so that a sense of two way human participation can be generated. Real time video conferencing is an example.

Educational Requirements

An Open Learner needs to be able to use any multimedia application at any time. However since open learning is often undertaken in centres, the use of audio and video require particular thought. Obviously several sets Sound Blaster driven speakers will disturb those learners working on a computer based self test! Similar considerations occur for users of multimedia in a class situation, such as a language teaching application. Not only will a number of students be performing similar activities at the same time on a network but the lecturer must decide whether to control the activities via the media of the computer. The use of multi-party desktop conferencing with the lecturer chairing the running of the conferencing session, showing selected parts of a video etc. is a case in point.

Distance learners or users of multimedia will also be capable of having same impact on a network as several students playing the computer game Doom at lunchtime. Additionally the co-ordination of a learning activity must also be done by a lecturer over the network. So the role of the chair in multi-party video conferencing is crucial.

Cost

In education the main costs visible to multimedia users to date have been the cost of the computer platform, the CD-ROM, and the software. Network costs are usually borne centrally within an institution, or by JANET nationally. The increased cost of providing sufficient network capacity and ability of new networks to charge on the basis of bandwidth used mean that individual users will increasing have to consider the costs of access when designing or using multimedia applications. Additionally information providers, electronic publishers, etc. will begin to incorporate charging mechanisms in their systems].

The educational user ideally needs the costs of use to be well defined, in advance so applications teaching journalism, which make repeated access to Reuters databases would not be viewed favourably by educational managers.

Cost benefit analysis of multimedia distance learning, and open learning proposals will increasingly become the norm.

ISSUES FOR USERS

Characteristics of Multimedia

Multimedia can be as simple as a few images with some accompanying text to a multimedia presentation using video clips, sound, images animation and text. Multimedia files to use a lot of data when in a digital format. Video is the most demanding. A PAL signal when digitised can require a data rate of 170 Mbps. Audio is less demanding but still requires 1.3 MByte for a 1 minute clip using a Sound Blaster Pro system at 22 kHz sampling rate. Still images require use more data proportional to their size. Synchronisation of sound and video is important. Sound is likely to breake up if parts of it are lost or delayed in storage or transmission. Video is less vulnerable to loss (depending on the application), but still requires all of the picture to be on the screen at the same time and is also vulnerable to jitter. Jitter could be controlled in some applications if the sender of the isochronous video data time stamps each piece of data when it is generated, using a universal time source, and then sends the data to the receiver. The receiver reads a piece of data in as soon as it is received and store it. The receiver processes each piece of data only at the time equal to the data's time stamp plus the maximum transit delay. Thus isochronicity of the video would be restored.

An example estimate of the requirements made by voice and video on an ATM network is given below.

Parameter	***Interactive Voice***	***Non-Interactive Video @ 30Mbps***
Delay	200 msec	1000 msec
Jitter	1 msec	5 msec
Throughput	8.8 Kbytes/sec	4.1 MBytes/sec
Average Throughput	3.9 kbytes/sec	4.1 Mbytes/sec
Packet sequencing required		Yes Yes
Absence of packet duplication		Yes Yes
Setup time	0.8 sec	15 sec

(Ferrari RFC 1193 Requirements for Real-Time Services November 1990).

To summarise, multimedia data is large, sensitive to delay and loss of data.

To accommodate these characteristics techniques used by the telecommunication networks to carry telephone and television traffic are required. These include compression of data, and methods of timing the transmission and replay of multimedia. Data networks and computers have been built in a different way (they are asynchronous) to telephone and TV networks (which are isochronous).

Compression

Compression algorithms and techniques are critical to the viability of multimedia networking. Uncompressed digital television requires about 140 Mbps. Since few users have this sort of network access compression is the only hope for the widespread deployment of digital video and multimedia. Compression techniques depend on algorithms implemented in software or hardware. The use of hardware is important still to enable rapid compression, and also speeds de-compression. At this time the cost of hardware is still high, from µ200 to µ350 for a MPEG video compression PC card. Sound cards can implement proprietary compression, and software only video compression is available in products like Microsoft Video for Windows, or for UNIX workstations.

While compression can ease the demands on networks and storage media there are several trade-offs. Since some compression techniques remove information considered to be less important a loss in resolution may result. Once material is compressed the algorithms may prevent access to single frames of video for viewing or editing. The cost of complex hardware and software and compression and decompression delay are other factors important to users.

Different uses require different compression methods. Video conferencing must be done in real time so fast encoding and decoding is needed. This is the aim of the H.261 standard. Video film distribution via cable networks, radio or CD is essentially a playback process, so encoding is not time critical, and decoding should be easy to implement to reduce consumer costs. The MPEG standards address these applications.

MIDI encoding of audio notes is not really a compression method, but almost another form of media. Inevitably, successful compression techniques encourage the design of applications which require higher bandwidths still, such as Super Definition TV which will also require appropriate compression.

COMPRESSION STANDARDS

Uncompressed PAL video as a digital signal needs 140 to 270 Mbps, while uncompressed digital HDTV needs 1.2 Gbps. A digitised colour picture at 35 mm film resolution needs about 80 Mbyte. One minute of 8 bit sound sampled at 22 kHz needs 1.3 Mbyte. These requirements are large for storage and transmission in computer networks. Fortunately many applications can live with video or sound that is not perfect. It is therefore possible to produce coding schemes that will drop certain information, or make a guess at likely values. Such coding can result in compression ratios of up to 200 times, so PAL video can be compressed to 1.5 Mbps. Compression must be done to a standard to enable decoding.

However their are many compression methods. Microsoft Video for Windows, Quicktime are proprietary examples. A standard called JPEG is used for still colour images. For video conferencing a standard called H.261 which is part of the H.320 standards family is used. Further video standards are MPEG-1 and MPEG-2 which gives a higher quality picture to MPEG-1, but at a higher data rate. MPEG is likely to be important for video on demand. An MPEG decoder card is available for about £300. Compression of video to MPEG can now be done in real time, but the hardware to support this is expensive.

JPEG COMPRESSION

The JPEG (Joint Picture Experts Group) [ISO10] is the first international digital compression standard for multi-level continuous tone still, black and white or colour images. It typically compresses images to 1/10 or 1/50 of their original size. It is based on use of a discrete cosine transform and requires same level of processing to compress and decompress an image. JPEG aims are :-

- To be applicable to any kind of continuous tone digital source image.
- To be able to be implemented in hardware or software at reasonable cost.
- To support the following modes of operation :
  - (a) Sequential encoding, *i.e.* left to right and top to bottom scanning.
  - (b) Progressive encoding, using multiple scans so that the image builds up gradually.
  - (c) Lossless encoding, in which the compresses image can be decompressed to be identical to the original.
  - (d) Hierarchical encoding in which the image is encoded at multiple resolutions, so that lower

resolution displays can be accessed without having to decompress the full resolution image.

Most implementations have been only of sequential encoding. A 10 MHz JPEG chip can typically compress a full page 24 bit colour 300 dpi image from 25 Mbyte to 1 Mbyte in about one second. JPEG takes each lock of 8 × 8 source image samples and codes them into coefficients. The most important coefficients are then preserved. JPEG compression and decompression by different systems are not guaranteed to be the same, but a an accuracy test is available.

JPEG has been used for full motion video by compressing each frame of the video. For a 640 × 480 pixel 24 bit colour JPEG compresses each frame of about 1 Mbyte by about a factor of 50, which results in a data rate of about 5 Mbits/sec for 30 frames/sec. For this high data rate most implementations use a small window in the PC of 256 × 240 pixels which reduces the data rate by a factor of five. Intel DVI compression does the same. Motion JPEG does not support audio compression which must be done separately.

MPEG COMPRESSION

The MPEG (Motion Picture Experts Group) [ISO11] has so far defined two compression algorithms.: MPEG-1 and MPEG-2. A common misconception is that MPEG-2 is a replacement for MPEG-1. Each algorithms has been specifically targeted at different bit rates. MPEG-2 runs at higher bit rates than MPEG-1. There are no firm constraints in either algorithm and it is possible to run MPEG-1 video at very high rates. MPEG requires more processing power to compress video than decompress, so it is ideal for video film distribution. MPEG-1 chips on the market provide about a 200:1 compression to yield VHS quality video at 1.2 to 1.5 Mbps The MPEG specifications allow manufacturers to implement different proprietary, but MPEG compliant algorithms. There is therefore no guarantee the the output quality of MPEG encoders will be the same. Basically the user pays for what they see. MPEG takes advantage of temporal redundancy in video pictures by specify three type of pictures:-

- Intra Pictures or I-pictures are coded using only the information present in the picture itself using cosine transforms. I-pictures use about two bits per coded pixel and are used about every two seconds.
- Predicted Pictures or P-pictures are coded with respect to the nearest previous P or I-picture and use forward prediction of the video picture content.
- Bi-directional or B-pictures that use a past and future picture as a reference. B-pictures provide the most compression and average out noise. Typically two B-pictures will separate a P-picture.

Pictures may not be sent in the order in which they are displayed, if reference pictures are needed for reconstruction. MPEG also provides for synchronisation of audio and video streams.

Display Order						
I	B	B	P	B	B	P
Video Stream Order						
I	P	B	B	P	B	B

Stream Versus Display Ordering

Motion compensation is a technique used to enhance the compression of P and B-pictures by examining the spatial difference between pixel blocks within the picture.

MPEG-2 uses many similar techniques, and allows for many different resolutions and frame rates. It also takes advantage of motion prediction between video fields, enabling higher compression ratios. Both MPEG-1 and MPEG-2 can run at reduced resolutions by reducing the number of pixels to be encoded by a factor of two before compression.

There are optimal choices for the use of MPEG-1 and MPEG-2. Below link rates of 3.5 Mbps MPEG-1 provides better video quality (VHS standard) if the pixel input is reduced. However for data rates above 5.0 Mbps MPEG-2 provides better quality, Super VHS or above video.

TRANSMISSION MEDIA

Copper Conductors

Copper conductors twisted together are the basis of the telephone network. As used in the analogue telephone network, with modems turning digital data into analogue tones, the data rate is limited to around 28 kbps. The introduction of the Integrated Services Digital Network (ISDN) led to the use of improved modulation and coding schemes.

In standard ISDN the user can access up to 128 kbps of data at distances up to 100 metres from an ISDN socket [BT94].

Local Area networks also use copper cable twisted pairs. Ethernet running at 10 Mbps can also operate at up to 100 metres or more, depending on the grade of the cable. These physical networks will be available and economic to operate for many years. For this reason considerable effort has gone into upgrading the network protocols that can be used on them for high bit rate real time multimedia.

A recent development is Asymmetric Digital Subscriber Lines (ADSL) [Bellcore92] technology which is aimed at using two wire copper loops at data rates of 1.544 Mbps in the network to user direction and about 600 kbps from the user to network. This is achieved by using better modulation and coding techniques. The driving force behind this technology is the delivery of on demand compressed VHS quality video to the home by Telecommunication operators wishing to compete with cable TV operators. If this technology becomes commercially available (it is undergoing trials) then the opportunity exists to deliver several ISDN B channels or even low speed ATM to the subscribers premises. In the short term this technology is unlikely to be useful to the academic community.

Coaxial Cable

Coaxial cable is still an effective medium to deliver multimedia. Cable TV networks use coaxial cable. BT uses coaxial cable to deliver some of its Kilostream services, and Local Area Ethernet networks can operate over coaxial cable to the 10BASE5 and 10BASE2 specification. In general coaxial cable enables longer distance transmission at higher data rates than twisted pair cable, but is more expensive. One video conferencing provider has taken advantage of the potential additional bit rate available over coaxial cable and designed a proprietary video conferencing system for use over a LAN (C-Phone).

Because of the higher cost of coaxial cable new LAN installations which are to be based on copper technology most often use twisted pairs.

Optical Fibre

Optical fibre transmission has been a strong enabling factor in the creation of a high capacity digital network on which networked multimedia applications will depend. On campus, the provision of optical fibre enable the creation of a high capacity backbones to Local Area Networks. However optical fibre provision to the staff desktop or student seat is still more of an exception than a rule due to the higher costs. In principal optical fibre can carry almost any type of traffic, at high data rates. It should be noted that there are different type of optical fibre, single mode, multi-mode, with different sizes, and attenuation characteristics. When upgrading from one technology to another, *e.g.* FDDI to ATM this may become important.

Radio Systems

Radio systems are likely to impact on multimedia in two ways.

Firstly radio technology operating at microwave frequencies is available to provide wireless local area networks. Use of this technology for delivery of real time multimedia material should be treated carefully, because radio links are susceptible to fading, interference, random delays etc. For non real time use this technology is likely to perform as well as current Ethernet LANs.

Secondly the introduction of digital mobile systems termed 'GSM' from a French acronym means that mobile users will have access to the same digital networks as fixed users. So in theory a mobile user will be able to connect to an ISDN users application, or an ATM users, with the only difference being one of speed.

At present the bit rate of the mobile link is about 9.6 kbps. This may increase. With improvement in compression techniques the possibility of video conferencing on the move arises. Whether academics will need to adapt their video conferencing equipment to communicate with students stuck the bus and late for a lecture is a matter for speculation!

LAN NETWORK TECHNOLOGIES

ETHERNET AT 10 MBPS

There is a vast installed base (about 40 million Ethernet nodes) of 10 Mbps Ethernet and 4 or 16 Mbps Token Ring

Local Area Networks (LANs) using coaxial cable or twisted copper wire.

Most new LANs used twisted copper wire cable. Ethernet uses a contention method to enable workstations attached to the same cable to share the data bandwidth on the cable.

Nodes transmit to the network on demand, but continually monitor the network to see if another node is transmitting at the same time. If this occurs both nodes cease transmission, and try again later at random intervals. The throughput with such a system is limited to about half the available bandwidth or 5 to 6 Mbps. Multimedia file servers can be attached to such LANs. Both Novell and Microsoft have developed software to support video on these traditional LANs.

Video in an Intel DVI or MPEG form requires 1 to 2 Mbps per user, so it is clear that only a handful of users can run video applications simultaneously. Audio requires lower bandwidth but is sensitive to unpredictable delays. For instance an intensive file transfer or video stream could prevent an audio application from transmitting onto the LAN for several tens of milliseconds, which is sufficient to reduce the intelligibility of voice. Improved performance is possible if LAN segments are divided up into segments with only a few attached users.

Switched Ethernet takes this approach to the limit by enabling only one user to access a single segment which is then connected to a higher capacity network at a central hub by a switch. The price per port of Ethernet switches ranges from £400 to £1000. The use of a switch permits the filtering of packets based on address. Switching times need to be fast enough to deliver packets at the basic 10 Mbps Ethernet line rate. An Ethernet switch would be equipped with a higher speed ATM or FDDI interface to other networks. However even with one user per Ethernet Switch port or segment contention between different applications on the same machine or incoming and outgoing traffic can occur.

FAST ETHERNET

The IEEE has set out the aims of the 100 Mbps Ethernet LAN standard [Rame93] as follows:-

- Line Rate of 100 Mbps
- 100 metre distances permissible to the Ethernet Hub
- Category 3 to 5 twisted pair operation
- Equivalent error rates to 10 Mbps
- Compliance with electromagnetic compatibility standards
- Simultaneous support of 10 Mbps and 100 Mbps
- Use of RJ-45 connector.

Most of the 10 Mbps parameters including the CSMA/CD media access protocol will remain unchanged and the standard is being named 100Base-T. Another 100 Mbps standard called 100VG-AnyLAN [LAN95] has also been proposed and changes the media access protocol to a demand priority system. Fast Ethernet products are only just starting to appear. Use of both standards will require a change out of workstation and hub cards. The cost of 100 Mbps Ethernet cards will be targeted to be comparable with high performance 10Base-T cards.

FDDI

The Fibre Distributed Data Interface (FDDI) [Minoli93] was the only standards based technology operating at 100 Mbps for some time. FDDI has experienced slow market penetration due to the high cost of cards, still around £600. The first version of FDDI was developed as a campus trunk network for data. Logically FDDI consists of a dual ring, but it may be implemented as a physical star. Key feature of FDDI include :-

- Shared medium based on a token passing medium access control.
- Compatibility with IEEE 802 LANs.
- Ability to use a wide range of physical medium including multi-mode fibre, single mode fibre, shield and unshielded twisted pair.
- Operation at 100 Mbps
- Support for 500 workstations
- Maximum fibre length of 200 km.
- Ability to allocate bandwidth dynamically so that both ‘synchronous and asynchronous services can be provided.

An upgraded FDDI standard called FDDI II has been designed. In addition to the data, packet switched mode in

FDDI an isochronous circuit switched service is made available by imposing a 125 micro second frame structure. The 100 Mbps bandwidth can be split between packet data and up to fifteen isochronous channel operating at 6.144 Mbps each.

WAN NETWORK SERVICES

FRAME RELAY

Frame relay is a connection oriented services operating at n x 64 kbps or 2.048 Mbps. It has evolved from X.25 packet switching and aims to reduce network delays, protocol overheads and equipment cost. Error correction is done on an end to end basis rather than a link to link basis as in X.25 switching. Frame relay can support multiple users over the same line and can establish a permanent virtual circuit or a switched virtual circuit.

Like ATM it is a protocol which must be carried over a physical link such as a Kilostream or Megastream link. While useful for connection of LANs, the combination of low throughput, delay variation and frame discard when the link is congested will limit its usefulness to multimedia.

SMDS

The Switched Multi-megabit Data Service (SMDS) [King93] is a new switched broadband data service. One of the first users of the service in the UK have been SuperJANET sites. SMDS provides a switched connectionless data service at speeds of 34 Mbps (at present) for connection of LANs. It uses variable length packets up to 9188 bytes in length, each of which carries an address in the E.164 format (ISDN uses this address format too) SMDS packets are transported in the public network using the Distributed Queue Dual Bus (DQDB) IEE 802.6 standard which uses packets fixed at 53 bytes. There is some overhead from this conversion process which reduces the bandwidth available to users to about 75% of the line speed. There are several access classes that limit the sustained data rate and burst data rate that can be injected into the network by a user. These access restrictions may result in discard of packets that exceed a certain limit. SMDS does not support timing. The higher speeds of SMDS will be of benefit to multimedia applications seeking to transfer large volumes of data quickly, but the lack of a time structure will reduce the video conferencing quality obtainable.

THE INTEGRATED SERVICES DIGITAL NETWORK (ISDN)

In the real world the delivery of multimedia requires a widespread network capable of delivering at high data rates. The current implementation of ISDN in the narrow band form is the best access and delivery medium available. ISDN is seen by many in the industry as the ramp through which multimedia networking will gain acceptance. The installed base of ISDN is growing rapidly (30,000 line per month in

Germany). ISDN is able to provide connections throughout the world. In Europe the Euro-ISDN agreements between operators is valuable.

ISDN offers point to point delivery, network access, and network interconnection for multimedia. Different data rates from 64 kbps up to 2 Mbps are commercially available which can meet many needs for transporting multimedia. Call set-up times are under one second.

ISDN will be the feeder network for broadband ISDN based on ATM standards. Initially the ISDN and ATM networks will be overlaid on top of each other, but users of ISDN will eventually be able to call an ATM user directly and be allocated an appropriate amount of bandwidth. The development of 'middleware' will enable applications to communicate over mixed networks.

Although ISDN could be cheaper, particularly in the UK, it is likely to be cheaper than ATM connections and more widespread in availability for a long time. It is therefore an important tool in bringing multimedia applications to a wide range of users. The idea that multimedia can only be delivered on broadband networks is erroneous as the assertion that only a Macintosh can deliver multimedia. The cost of ISDN hardware was high, but is now decreasing. Terminal adapters are available from £400 upwards, and PC cards for £300 upwards. Video conferencing cards cost around £3000, (BT's VC8000 card). Costs of ISDN equipment are much lower in

Germany and some of these products are beginning to appear in the UK under the Euro-ISDN banner.

British Telecom are pursuing a strategy to make ISDN the preferred option for all multiple (2 or more) exchange line requests by the mid 1990s. ISDN is accessed through one of two services, named by the CCITT as Basic Rate Access (BRA) and Primary Rate Access (PRA).

Basic Rate Access (BRA) provides an ISDN user with simultaneous access to two 64 kbps data channels using the existing twisted pair copper telephone cable. The B.T. basic rate ISDN service is called ISDN2. The connection cost of ISDN2 is currently £300. Rental for the equivalent of two PSTN telephone lines is £384 per year.

Each data channel is referred to as a B-channel and can carry voice or data. Another channel, the D-channel, operates at 16 kbps and is used for signalling between user devices and the ISDN. The total data rate of BRA is therefore 144 kbps. The two B-channels and the single signalling channel give rise to the term '2B+D'. BRA is also referred to as I.420, after the CCITT recommendation. Basic rate ISDN is intended for low capacity usage, such as that required for small businesses.

British Telecom's primary rate ISDN service is known as ISDN30. This service is generally available throughout the UK and is based on the CCITT recommendations for primary rate ISDN. Mercury Communications Limited also offer a primary rate service known as 2100 Premier. Although this service is largely based on CCITT recommendations, it still utilises the some proprietary signalling.

Primary rate access can carry 30 independent voice or data channels, each at 64 kbps. The structure has a 64 kbps D-channel for signalling between devices and the network, and a 64 kbps channel for synchronisation and monitoring. The total data rate of PRA is 2.048 Mbps.

Primary rate access is often referred to as '30B+D' because of the number of B-channels and D-channels, or I.421 because of the CCITT recommendation from which it is taken. This form of access is primarily intended for use in situations which require a large transmission capacity, such as when organisations make voice and data calls through an Integrated Services PBX.

There are two standard ISDN connectors. For accessing basic rate ISDN, an RJ-45 type plug and socket (similar to a telephone plug) is used using unshielded twisted pair cable. Access to primary rate ISDN is through a coaxial cable.

The ISDN passive bus, which can be a maximum of 1 km in length, is a cable which in user premises. It enables up to eight user devices to be attached to the basic rate ISDN interface. Since there are only two B-channels, only two of the eight devices can communicate at any one time. For this reason, each device must contend for access to the passive bus.

ISDN signalling information, carried in the D-channel, is used to establish, monitor and control ISDN connections between users as well as instigating, the audible ringing or engaged tones. The ISDN numbering system is similar to the contemporary telephone numbering system. Each B-channel has its own unique directory number which allows access to different terminal types (such as telex or facsimile devices). Each terminal type has an identity code which ensures that it only communicates with similar terminals.

The equipment available for ISDN includes Terminal Adapters, ISDN internal computer Terminal Adapter cards, Video Conferencing PC cards, and LAN access gateways or bridges, some of which are based on PC cards or stand alone boxes. Products are available from in this country from the USA, UK, France, and Germany. The market for ISDN is most developed in Germany.

Internal Terminal Adapters from Germany will all inter-work with each other, products developed in the UK are all totally and individually proprietary and will not inter-work in many cases. Many manufacturers are awaiting the dust to settle on the competing application programming interface standards from the European PTT body ETSI.

It is possible to avoid all the problems of API standards for internal computer adapters by using an external ISDN Terminal Adapter. Since the speed of most serial ports on a PC has been limited to about 19.2 kbps until recently, this approach has not been viable. However recently internal PC cards which will work asynchronously up to 115 kbps have appeared, which could have applications in multimedia work when used with an appropriate external Terminal Adapter.

2

Multimedia and Internet Protocols

EXISTING INTERNET PROTOCOLS

The Internet is a more of a phenomena than a network, but is important when discussing multimedia because a popular Internet Application, the World Wide Web is capable of accessing and displaying multimedia formats such as pictures, audio and video. The current Internet has thrived and grown due to the existence of TCP implementations for a wide variety of classes of host computers. These various TCP implementations achieve robust inter-operability by a “least common denominator” approach to features and options.

The system of connected networks which comprise the Internet has also been used to carry live audio and video.

Extensions to the TCP/IP protocols currently used have been proposed as Real Time Protocols (RTP). Broadcast of audio and video has taken place on the Multicast Backbone (MBONE), by allocating higher priority to audio and video information from within routers.

The MBONE is being developed as a technology for low cost multimedia. Multicasting within MBONE enables multiple destinations to share the same information without replication. Internet routers and workstation software require some modifications to support multicasting. A virtual network has been implemented over the IP network to bypass routers which do not support multicasting, and to enable some bandwidth to be reserved for multicasting. However audio and video on the MBONE must still compete with other traffic on parts of the network. This limits the quality of both the voice and video obtainable.

However current transport protocols exhibit some severe problems for high performance, especially for using hardware support. Existing protocols require a processing overhead which takes longer than the transmission time on high speed networks. For example, TCP places the checksum in the packet header, forcing the packet to be formed and read fully before transmission begins. ISO TP4 is even worse, locating the checksum in a variable portion of the header at an indeterminate offset, making hardware implementation extremely difficult.

Special purpose transport protocols have been developed. Examples include special purpose transport protocols such as UDP (user datagram protocol), RDP (reliable datagram protocol), NVP (network voice protocol), PVP (packet video protocol) and XTP (Xpress Transfer Protocol), XTP fixes header and trailer sizes to simplify processing and places error correction in the trailer so that the code can be calculated while information bits are being transmitted. Flow, error and rate control are also modified in XTP. Examples of XTP applications include :-
A video-mail demo over XTP/FDDI that uses a proprietary Fluent multimedia interface and standard JPEG compression. This PC-based demo delivers full frame, full colour, 30 frames/s video from any network disk to a remote VGA screen. Voice can be multicasted over XTP/FDDI. A simple multicast is distributed to a group with a latency of around 25 ms, where the latency represents delay from the voice signal from the microphone to the audio signal to the speaker. Commercially, Starlight Networks Inc., migrated a subset of XTP into the transport layer of its video application server. By using XTP rate control, full-motion, full-screen compressed video is delivered at a constant 1.2 Mbps, over switched-hub Ethernet to work stations. This network delivers at least 10 simultaneous video streams. The Internet physically depends on the capabilities of the underlying networks. If TCP/IP protocols are to be used in a world

equipped with ATM capable of transporting audio and video efficiently then any adaptation of current TCP/IP protocols will need to be tailored to the needs of multimedia.

New Internet Protocols

A successor to the current version of TCP/IP Version 4 is being discussed. While some of motivation behind this is due to the need to increase the address space available, the opportunity is also being taken to review the need to increase the performance of Internet protocols, for multimedia applications. The Internet Engineering Task Force has resolved to move towards develop a replacement for the current TCP/IP Version 4 called IPng (IP next generation) [Brazdziunas94]. Effective support for high quality video and audio streams is one of the critical capabilities that is being called for to capture the attention of network operators and information providers of interactive broadband services (*e.g.*, cable television industry and partners). Such additional features will also help overcome resistance to change. The intention is that IPng should last for the next 20 years.

The delivery of digital video and audio programs requires the capability to do broadcasting and selective multicasting efficiently. The interactive applications that the future cable networks will provide will be based on multimedia information streams that will have real time constraints. The largest fraction of the future broadband traffic will be due to real time voice and video streams. It will be necessary to provide performance bounds for bandwidth, jitter, latency and loss parameters, as well as synchronisation between media streams related by an application in a given session.

The potential for IPng to provide a universal internetworking solution is a very attractive possibility, but there are many hurdles to be overcome. One of these is that a new deployment of IPng threatens the existing network investments that business has made and the other is that business users actually buy applications — not networking technologies. Some of the the aims of IPng development relevant to multimedia are set out below:-

Two aspects are worth mentioning. First, the quality of service parameters are not known ahead of time, and hence the network will have to include flexible capabilities for defining these parameters. For instance, MPEG-2 packetised video might have to be described differently than G.721 PCM packetised voice, although both data streams are real time traffic channels.

Network media speeds are constantly increasing. It is essential that the Internet switching elements (routers) be able to keep up with the media speeds. A proper IPng router should be capable of routing IPng traffic over links at speeds that are capable of fully utilising an ATM switch on the link.

Processing of the IPng header, and subsequent headers (such as the transport header), can be made more efficient by aligning fields on their natural boundaries and making header lengths integral multiples of typical word lengths (32, 64, and 128 bits have been suggested) in order to preserve alignment in following headers. Optimising the header's fields and lengths only for today's processors may not be sufficient for the long term. Processor word and cache-line lengths, and memory widths are constantly increasing.

There are now many different LAN, MAN, and WAN media, with individual link speeds ranging from a ones-of-bits per second to hundreds of gigabits per second. There will be multiple-access and point-to-point links on a switched and permanent basis. At a minimum, media running at 500 gigabits per second will be commonly available within 10 years. Switched circuits include both "permanent" connections such as X.25 and Frame Relay services and "temporary" types of dial up connections similar to today's SLIP and dial up PPP services, and perhaps, ATM SVCs. Any IPng will need to operate over ATM. However, IPng still must be able to operate over other, more "traditional" network media. A host on an ATM network must be able to inter-operate with a host on another, non-ATM, medium.

Multicasting has been used with a limited degree of success to support audio and video broadcasts. Tests at ULC used DVI video compression with a data rate of up to 600 kbps and achieved a frame rate of up to 5 frames per second. Tests of H.261 video, also from ULC encountered delays of up to 12 seconds on the IP network. Some of this delay could be buffered out, raising the average delay. The conclusions were that slow TCP error recovery mechanism was inappropriate, and the UDP protocol may give better results.

On mixed protocol networks IPv4 currently uses the local media broadcast address to multicast to all IP hosts. This method is detrimental to other protocol traffic on a network. The ability to restrict the range of a multicast to specific networks is also important. Currently, large-scale multicasts are routed manually through the Internet. User configurable Multicast Addressing is vital to support future applications such as remote conferencing. For many reasons, such as accounting, security and multimedia, it is desirable to treat different packets differently in the network. For example, multimedia is now on our desktop and will be an essential part of future networking. Multimedia applications need to acquire differing grades of network service, for voice, video, file transfer, etc. It is essential that this service information be propagated around the network. To support multimedia features will be needed such as policy-based routing, flows, resource reservation, type-of-service and quality-of-service.

NETWORK CABLING

Cable is the medium through which information usually moves from one network device to another. There are several types of cable which are commonly used with LANs. In some cases, a network will utilize only one type of cable, other networks will use a variety of cable types.

The type of cable chosen for a network is related to the network's topology, protocol, and size. Understanding the characteristics of different types of cable and how they relate to other aspects of a network is necessary for the development of a successful network. The following sections discuss the types of cables used in networks and other related topics.

- Unshielded Twisted Pair (UTP) Cable
- Shielded Twisted Pair (STP) Cable
- Coaxial Cable
- Fibre Optic Cable
- Cable Installation Guides
- Wireless LANs
- Unshielded Twisted Pair (UTP) Cable.

Twisted pair cabling comes in two varieties: shielded and unshielded. Unshielded twisted pair (UTP) is the most popular and is generally the best option for school networks.

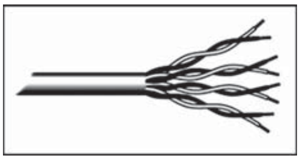

Fig. Unshielded Twisted Pair.

Categories of Unshielded Twisted Pair

Category	***Speed***	***Use***
1	1 Mbps	Voice Only (Telephone Wire)
2	4 Mbps	LocalTalk & Telephone (Rarely used)
3	16 Mbps	10BaseT Ethernet
4	20 Mbps	Token Ring (Rarely used)
5	100 Mbps (2 pair)	100BaseT Ethernet
	1000 Mbps (4 pair)	Gigabit Ethernet
5e	1,000 Mbps	Gigabit Ethernet
6	10,000 Mbps	Gigabit Ethernet

The quality of UTP may vary from telephone-grade wire to extremely high-speed cable. The cable has four pairs of wires inside the jacket. Each pair is twisted with a different number of twists per inch to help eliminate interference from adjacent pairs and other electrical devices. The tighter the twisting, the higher the supported transmission rate and the greater the cost per foot. The EIA/TIA (Electronic Industry Association/Telecommunication Industry Association) has established standards of UTP and rated six categories of wire (additional categories are emerging).

UNSHIELDED TWISTED PAIR CONNECTOR

The standard connector for unshielded twisted pair cabling is an RJ-45 connector. This is a plastic connector that looks like a large telephone-style connector. A slot allows the RJ-45 to be inserted only one way. RJ stands for Registered Jack, implying that the connector follows a standard borrowed from the telephone industry. This standard designates which wire goes with each pin inside the connector.

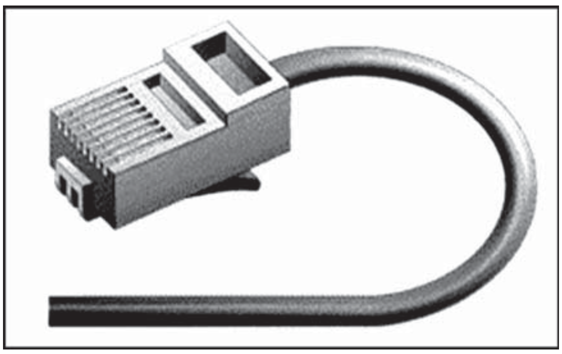

Fig. RJ-45 Connector.

SHIELDED TWISTED PAIR (STP) CABLE

Although UTP cable is the least expensive cable, it may be susceptible to radio and electrical frequency interference (it should not be too close to electric motors, fluorescent lights, etc.). If you must place cable in environments with lots of potential interference, or if you must place cable in extremely sensitive environments that may be susceptible to the electrical current in the UTP, shielded twisted pair may be the solution. Shielded cables can also help to extend the maximum distance of the cables.

Shielded twisted pair cable is available in three different configurations:

- Each pair of wires is individually shielded with foil.
- There is a foil or braid shield inside the jacket covering all wires (as a group).
- There is a shield around each individual pair, as well as around the entire group of wires (referred to as double shield twisted pair).

COAXIAL CABLE

Coaxial cabling has a single copper conductor at its centre. A plastic layer provides insulation between the centre conductor and a braided metal shield. The metal shield helps to block any outside interference from fluorescent lights, motors, and other computers.

Although coaxial cabling is difficult to install, it is highly resistant to signal interference. In addition, it can support greater cable lengths between network devices than twisted pair cable. The two types of coaxial cabling are thick coaxial and thin coaxial.

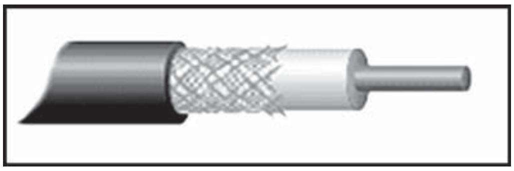

Fig. Coaxial Cable.

Thin coaxial cable is also referred to as thinnet. 10Base2 refers to the specifications for thin coaxial cable carrying Ethernet signals. The 2 refers to the approximate maximum segment length being 200 meters. In actual fact the maximum segment length is 185 meters. Thin coaxial cable has been popular in school networks, especially linear bus networks.

Thick coaxial cable is also referred to as thicknet. 10Base5 refers to the specifications for thick coaxial cable carrying Ethernet signals. The 5 refers to the maximum segment length being 500 meters. Thick coaxial cable has an extra protective plastic cover that helps keep moisture away from the centre conductor. This makes thick coaxial a great choice when running longer lengths in a linear bus network. One disadvantage of thick coaxial is that it does not bend easily and is difficult to install.

COAXIAL CABLE CONNECTORS

The most common type of connector used with coaxial cables is the Bayone-Neill-Concelman (BNC) connector. Different types of adapters are available for BNC connectors, including a T-connector, barrel connector, and terminator. Connectors on the cable are the weakest points in any network. To help avoid problems with your network, always use the BNC connectors that crimp, rather screw, onto the cable.

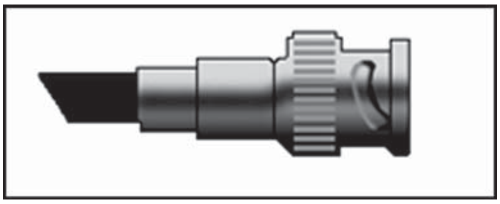

Fig. BNC Connector.

FIBRE OPTIC CABLE

Fibre optic cabling consists of a centre glass core surrounded by several layers of protective materials. It transmits light rather than electronic signals eliminating the problem of electrical interference. This makes it ideal for certain environments that contain a large amount of electrical interference. It has also made it the standard for connecting networks between buildings, due to its immunity to the effects of moisture and lighting.

Fibre optic cable has the ability to transmit signals over much longer distances than coaxial and twisted pair. It also has the capability to carry information at vastly greater speeds. This capacity broadens communication possibilities to include services such as video conferencing and interactive services. The cost of fibre optic cabling is comparable to copper cabling; however, it is more difficult to install and modify. 10BaseF refers to the specifications for fibre optic cable carrying Ethernet signals.

The centre core of fibre cables is made from glass or plastic fibres. A plastic coating then cushions the fibre centre, and kevlar fibres help to strengthen the cables and prevent breakage. The outer insulating jacket made of teflon or PVC.

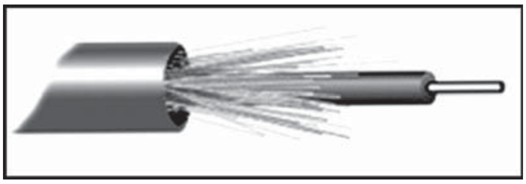

Fig. Fibre Optic Cable.

There are two common types of fibre cables — single mode and multimode. Multimode cable has a larger diameter; however, both cables provide high bandwidth at high speeds. Single mode can provide more distance, but it is more expensive.

Specification	***Cable Type***
10BaseT	Unshielded Twisted Pair
10Base2	Thin Coaxial
10Base5	Thick Coaxial
100BaseT	Unshielded Twisted Pair
100BaseFX	Fibre Optic
100BaseBX	Single mode Fibre
100BaseSX	Multimode Fibre
1000BaseT	Unshielded Twisted Pair
1000BaseFX	Fibre Optic
1000BaseBX	Single mode Fibre
1000BaseSX	Multimode Fibre

INSTALLING CABLE-SOME GUIDELINES

When running cable, it is best to follow a few simple rules:

- Always use more cable than you need. Leave plenty of slack.

- Test every part of a network as you install it. Even if it is brand new, it may have problems that will be difficult to isolate later.
- Stay at least 3 feet away from fluorescent light boxes and other sources of electrical interference.
- If it is necessary to run cable across the floor, cover the cable with cable protectors.
- Label both ends of each cable.
- Use cable ties (not tape) to keep cables in the same location together.

WIRELESS LANS

More and more networks are operating without cables, in the wireless mode. Wireless LANs use high frequency radio signals, infrared light beams, or lasers to communicate between the workstations, servers, or hubs. Each workstation and file server on a wireless network has some sort of transceiver/antenna to send and receive the data. Information is relayed between transceivers as if they were physically connected. For longer distance, wireless communications can also take place through cellular telephone technology, microwave transmission, or by satellite.

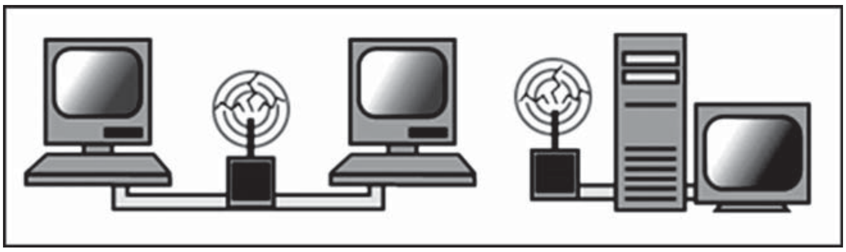

Wireless networks are great for allowing laptop computers, portable devices, or remote computers to connect to the LAN. Wireless networks are also beneficial in older buildings where it may be difficult or impossible to install cables.

The two most common types of infrared communications used in schools are line-of-sight and scattered broadcast. Line-of-sight communication means that there must be an unblocked direct line between the workstation and the transceiver.

If a person walks within the line-of-sight while there is a transmission, the information would need to be sent again. This kind of obstruction can slow down the wireless network. Scattered infrared communication is a broadcast of infrared transmissions sent out in multiple directions that bounces off walls and ceilings until it eventually hits the receiver. Networking communications with laser are virtually the same as line-of-sight infrared networks.

WIRELESS STANDARDS AND SPEEDS

The Wi-Fi Alliance is a global, non-profit organization that helps to ensure standards and interoperability for wireless networks, and wireless networks are often referred to as Wi-Fi (Wireless Fidelity).

The original Wi-Fi standard (IEEE 802.11) was adopted in 1997. Since then many variations have emerged (and will continue to emerge). Wi-Fi networks use the Ethernet protocol.

Standard	*Max Speed*	*Typical Range*
802.11a	54 Mbps	150 feet
802.11b	11 Mbps	300 feet
802.11g	54 Mbps	300 feet
802.11n	100 Mbps	300+ feet

WIRELESS SECURITY

Wireless networks are much more susceptible to unauthorized use than cabled networks. Wireless network devices use radio waves to communicate with each other. The greatest vulnerability to the network is that rogue machines can "eves-drop" on the radio wave communications. Unencrypted information transmitted can be monitored by a third-party, which, with the right tools (free to download), could quickly gain access to your entire network, steal valuable passwords to local servers and online services, alter or destroy data, and/or access personal and confidential information stored in your network servers. To minimize the possibility of this, all modern access points and devices have configuration options to encrypt transmissions. These encryption methodologies are still evolving, as are the tools used by malicious hackers, so always use the strongest encryption available in your access point and connecting devices.

A NOTE ON ENCRYPTION: As of this writing WEP (Wired Equivalent Privacy) encryption can be easily hacked with readily-available free tools which circulate the internet. WPA and WPA2 (WiFi Protected Access versions 1 and 2) are much better at protecting information, but using weak passwords or passphrases when enabling these encryptions may allow them to be easily hacked. If your network is running WEP, you must be very careful about your use of sensitive passwords or other data. Three basic techniques are used to protect networks from unauthorized wireless use. Use any and all of these techniques when setting up your wireless access points:

Encryption

Enable the strongest encryption supported by the devices you will be connecting to the network. Use strong passwords (strong passwords are generally defined as passwords containing symbols, numbers, and mixed case letters, at least 14 characters long).

Isolation

Use a wireless router that places all wireless connections on a subnet independent of the primary private network. This protects your private network data from pass-through internet traffic.

Hidden SSID

Every access point has a Service Set IDentifier (SSID) that by default is broadcast to client devices so that the access point can be found. By disabling this feature, standard client connection software won't be able to "see" the access point. However, the eves-dropping programs discussed previously can easily find these access points, so this alone does little more than keep the access point name out of sight for casual wireless users.

Advantages of wireless networks:

- Mobility- With a laptop computer or mobile device, access can be available throughout a school, at the mall, on an airplane, etc. More and more businesses are also offering free WiFi access ("Hot spots").
- Fast setup- If your computer has a wireless adapter, locating a wireless network can be as simple as clicking "Connect to a Network" — in some cases, you will connect automatically to networks within range.
- Cost- Setting up a wireless network can be much more cost effective than buying and installing cables.
- Expandability- Adding new computers to a wireless network is as easy as turning the computer on (as long as you do not exceed the maximum number of devices).

Disadvantages of Wireless Networks:

- Security- Be careful. Be vigilant. Protect your sensitive data with backups, isolated private networks, strong encryption and passwords, and monitor network access traffic to and from your wireless network.
- Interference- Because wireless networks use radio signals and similar techniques for transmission, they are susceptible to interference from lights and electronic devices.
- Inconsistent connections- How many times have you hears "Wait a minute, I just lost my connection?"

Because of the interference caused by electrical devices and/or items blocking the path of transmission, wireless connections are not nearly as stable as those through a dedicated cable.

- Speed- The transmission speed of wireless networks is improving; however, faster options (such as gigabit Ethernet) are available via cables. If you are only using wireless for internet access, the actual internet connection for your home or school is generally slower than the wireless network devices, so that connection is the bottleneck. If you are also moving large amounts of data around a private network, a cabled connection will enable that work to proceed much faster.

3

Computer Networks

A computer network is two or more computers connected together using a telecommunication system for the purpose of communication and sharing resources". Ask any computer network expert to simplify this definition to you and you will start a debate on how it should not be just two computers but three. Simply put a network is a means of communication between computers. Within a given network, computers can send files, e-mails and other correspondence to each other. Even things like instant messaging, is set up within a computer's network. There are two basic types of computer Networks. Lan and Wan.

LAN

LAN or Local Area Network is the most common kind of network set up. There are two ways to connect a LAN network.

The simplest and easiest way is the peer-to-peer connection network. This is when two or more computers are directly connected to each other. For example if there were four computers in the network, computer 1 would be connected to computer 2, computer 2 would be connected to computer 3 and computer 3 would be connected to computer 4.

This means each computer is dependent on the other. And if there were a network problem with any one computer, all of them would be affected. The other type if the client server connection. This is the type of connection where all the computers in a given network are connected to one central computer. This is a more complicated network but one that is much more efficient that peer-to-peer.

A local area network, or LAN, is a network of connected computers in a room, building, or set of buildings. Local area networks have been around since the beginning of computer use. A LAN is defined as a user network whereby data is sent at high rates between people located relatively close to each other. LANs do not usually make use of leased communication lines, but only means of communication that are provided by the installer of the network.

The Internet is a wide area network, or WAN, which is distinct from a LAN. In contrast to the term *Internet*, local area networks are often called *intranets*, though sometimes this term refers to a cluster of LANs associated with a particular company or organization but not connected to the larger Internet. A local area network uses a hub or router to connect computers together.

The means of communication is the omnipresent Ethernet cable or wireless wi-fi technology. These technologies offer data transfer rates running between 10 to 10000 Mbit/s. Larger, more important LANs have redundant lines or other backup protocols. In networked computers, the most popular communication protocol is TCP/IP. Smaller LANs may be temporary and used between friends to play computer games over the network.

Over a network, users can share files, view files, make changes to data on other computers if permitted, play movies or music on multiple computers at once, chat with instant messaging, send e-mails to each other, play games, and so on. All the advantages of the Internet apply, although they only include others on the LAN, and the data transfer rates are high.

Perhaps the most frequently employed use of a LAN is to connect users to the Internet with only one connected router. In modern times, we use broadband cable or DSL modems to connect to the Internet, and it would be clumsy to have a modem associated with every computer, so we simply plug the modem into a router and link the router to computers with Ethernet cables. Configuring a LAN can be intimidating at first, but contemporary operating systems have programmes that do most of the necessary configurations automatically, so setting up a local area network is pretty easy. A local area network (LAN) supplies networking capability to a group of computers in close proximity to each other such as in an office building, a school, or a home. A LAN is useful for sharing resources like files, printers, games or other applications. A LAN in turn often connects to other LANs, and to the Internet or other WAN.

Most local area networks are built with relatively inexpensive hardware such as Ethernet cables, network adapters, and hubs. Wireless LAN and other more advanced LAN hardware options also exist.

Examples

The most common type of local area network is an Ethernet LAN. The smallest home LAN can have exactly two computers; a large LAN can accommodate many thousands of computers. Many LANs are divided into logical groups called subnets. An Internet Protocol (IP) “Class A” LAN can in theory accommodate more than 16 million devices organized into subnets.

High Performance Network

The HPN is used for the exchange of large amounts of operational data. Two Force10 E600 Routers, interconnected via 4-way 10-Gigabit Ethernet aggregated links, provide connectivity between the High Performance Computing Facility (HPCF) and the Data-Handling System (DHS). The

HPCF network nodes are connected via 10-Gigabit Ethernet and all DHS nodes via Gigabit Ethernet aggregated links.

General Purpose Network

The GPN is used for all other traffic. It has at its core two Foundry BigIron RX-16 routers and at the edge seven Foundry Super-X switches. The core routers are interconnected via 4-way 10-Gigabit Ethernet aggregated links and have multiple Gigabit Ethernet uplinks to the edge routers. The core also includes two further Super-X switches that are dual-attached to the RX-16s via 10-Gigabit Ethernet.

The GPN provides connectivity to:

- The HPCF, the DHS and additional servers via Gigabit Ethernet ports in the core.
- The user desktops and laptops via Gigabit Ethernet ports on the edge switches.
- The firewalls (for the Wide Area Network and the Demilitarized Zone (DMZ) via Gigabit Ethernet ports in the core. The DMZ includes ECaccess, web servers, the mail gateway and DNS (Domain Name Servers).

The Hardware

Both the Force10 E600 chassis are populated with 24 10-Gigabit-Ethernet and 144 Gigabit Ethernet ports. For resiliency there are four power supplies, two CPU modules and nine switching fabric modules. Both the RX-16 chassis are populated with 8 10-Gigabit Ethernet ports and 144 Gigabit Ethernet ports. For resiliency there are seven power supplies, two CPU modules and four switching fabric modules. The Super-X chassis each contain up to 156 Gigabit Ethernet ports and dual power-supplies.

WAN

Definition

The wide area network, often referred to as a WAN, is a communications network that makes use of existing technology to connect local computer networks into a larger working network that may cover both national and international locations. This is in contrast to both the local area network and the metropolitan area network, which provides communication within a restricted geographic area. Here is how the wide area network functions, and why it is so important to communications today.

The concept of linking one computer network with another is often desirable, especially for businesses that operate a number of facilities. Beginning with the local area network and going up to the wide area network, this is most easily accomplished by using existing telephony technology. Essentially, fibre optics are used to create the link between networks located in different facilities.

Often, this means using standard phone lines, referred to as POTS, or employing PSTN (public switched telephone network) technology. During the 1990s, a third option, that of ISDN (integrated services digital network) solutions for creation a wide area network gained a great deal of popularity, mainly because the concept made it more cost effective to extend the network beyond national boundaries.

With coverage in a broad area, a wide area network allows companies to make use of common resources in order to operate. For example, many retail drugstores make use of a wide area network as part of their support to customers who fill prescriptions with one of their stores. Once in the common customer database for the pharmacy, the client is free to fill a prescription at any of the company's locations, even while vacationing in another state.

Companies also make good use of the wide area network as well. Internal functions such as sales, production and development, marketing and accounting can also be shared with authorized locations through this sort of broad area network application.

The concept of a wide area network as a means of taking individual location based computer networks and using them to create a unified computer network for the entire corporation means that employees can work from just about anywhere. Should one facility be damaged or rendered inaccessible due to natural disaster, employees simply move to another location where they can access the unified network, and keep on working.

Overview

WAN or Wide Area Network is when several LANs or independent computers are connected to a single, wider network. The Internet is the perfect example of WAN. E-mails, Chat Rooms and IMs all connect to the WAN of the Internet. WAN is much more complex and requires connecting devices or hubs from all over the world. The term Wide Area Network (WAN) usually refers to a network which covers a large geographical area, and use communications circuits to connect the intermediate nodes.

A major factor impacting WAN design and performance is a requirement that they lease communications circuits from telephone companies or other communications carriers. Transmission rates are typically 2 Mbps, 34 Mbps, 45 Mbps, 155 Mbps, 625 Mbps (or sometimes considerably more). Numerous WANs have been constructed, including public packet networks, large corporate networks, military networks, banking networks, stock brokerage networks, and airline reservation networks.

Some WANs are very extensive, spanning the globe, but most do not provide true global coverage. Organizations supporting WANs using the Internet Protocol are known as Network Service Providers (NSPs). These form the core of the Internet. Wide Area Networks, or WANs, connect a geographically diverse group of computers within a state, country, or even across several states or countries. WANs typically are connected by telephone lines, other types of communication lines, or radio waves.

Quite often, smaller local area networks (LANs) are linked together to form a WAN. This is accomplished via dedicated private lines, leased from telecommunications firms like Sprint and ATandT, or by Switched Multi-Megabit Data Services (SMDS) technology, developed in 1995 to eliminate the need for a leased line. WAN technology has been refined over a period of several decades. It first emerged in the mid-twentieth century with the advent of networks like ARPAnet. Developed in 1969 by the Department of Defence, ARPAnet and several other networks eventually evolved into the Internet, the largest WAN in the world.

The packet switching technology most commonly used with WANs surfaced in the 1960s, and standard packet switching protocol, known as X.25, was developed in 1976. To increase network speed, packet switching allows for the parceling of data into smaller chunks, known as packets, prior to transmission. These packets can travel independently via alternate routes, and they are reassembled once they reach their target.

Although X.25 remained the most popular WAN packet switching protocol for years, other packet switching protocols used with increasing frequency by WAN developers and administrators include the Internet standard, Transmission Control Protocol/Internet Protocol (TCP/IP), and Frame

Relay, used most often by WANs connected via high speed T-1 and T-3 lines.

WANs are used for a variety of purposes. A corporation with offices in several locations may use a WAN to form an intranet. Quite often, the individual offices will use their own LANs for things like internal messaging, data processing functions, and hardware and software sharing. When these LANs are joined together to form a WAN, similar data sharing and messaging capabilities become possible across a much broader geographic area.

Businesses wanting to link up with their suppliers or distributors may create a WAN as a means of establishing an extranet. For example, an extranet could provide a sales representative with electronic access to information in about the time it might take to deliver a product, or the availability of a product.

Some WANs bring together various types of communications, such as data, video, and voice. Some organizations, including companies, universities, research centres, hospitals, and libraries, use WANs to connect to the Internet. By connecting the NSP WANs together using links at Internet Packet Interchanges (sometimes called "peering points") a global communication infrastructure is formed.

NSPs do not generally handle individual customer accounts (except for the major corporate customers), but instead deal with intermediate organizations whom they can charge for high capacity communications. They generally have an agreement to exchange certain volumes of data at a certain "quality of service" with other NSPs.

So practically any NSP can reach any other NSP, but may require the use of one or more other NSP networks to reach the required destination. NSPs vary in terms of the transit delay, transmission rate, and connectivity offered.

The characteristics of the transmission facilities lead to an emphasis on efficiency of communications techniques in the design of WANs.

Controlling the volume of traffic and avoiding excessive delays is important. Since the topologies of WANs are likely to be more complex than those of LANs, routing algorithms also receive more emphasis.

Many WANs also implement sophisticated monitoring procedures to account for which users consume the network resources. This is, in some cases, used to generate billing information to charge individual users.

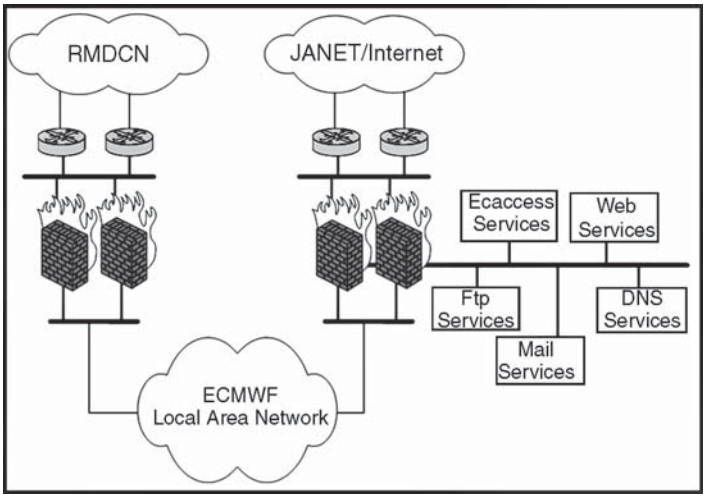

The size of a network is limited due to size and distance constraints. However networks may be connected over a high-speed communications link (called a WAN link) to link them together and thus becomes a WAN. WAN links are usually:

- Dial up connection
- Dedicated connection-It is a permanent full time connection. When a dedicated connection is used, the cable is leased rather than a part of the cable bandwidth and the user has exclusive use.
- Switched network-Several users share the same line or the bandwidth of the line.

  There are two types of switched networks:

  1. *Circuit switching*: This is a temporary connection between two points such as dial-up or ISDN.
  2. *Packet switching*: This is a connection between multiple points. It breaks data down into small packets to be sent across the network. A virtual circuit can improve performance by establishing

a set path for data transmission. This will shave some overhead of a packet switching network. A variant of packet switching is called cell-switching where the data is broken into small cells with a fixed length.

Connection Technologies

- X.25-This is a set of protocols developed by the CCITT/ITU which specifies how to connect computer devices over an internet work. These protocols use a great deal of error checking for use over unreliable telephone lines. They establish a virtual communication circuit. It uses a store and forward method which can cause about a half second delay in data reception when two way communications are used. Their speed is about 64Kbps. Normally X.25 is used on packed switching PDNs (Public Data Networks). A line must be leased from the LAN to a PDN to connect to an X.25 network. A PAD (packet assembler/disassembler) or an X.25 interface is used on a computer to connect to the X.25 network. CCITT is an abbreviation for International Telegraph and Telephone Consultative Committee. The ITU is the International Telecommunication Union.
- Frame Relay-devices at both sides of the connection handle Error checking. Frame relay uses frames of varying length and it operates at the data link layer

of the OSI model. A permanent virtual circuit (PVC) is established between two points on the network. Frame relay speed is between 56Kbps and 1.544Mbps. Frame relay networks provide a high-speed connection up to 1.544Mbps using variable-length packet switching over digital fibre-optic media. Frame relay does not store data and has less error checking than X.25.

- Switched Multi-megabit Data Service (SMDS)-Uses fixed length cell switching and runs at speeds of 1.533 to 45Mbps. It provides no error checking and assumes devices at both ends provide error checking.
- Telephone connections
  - Dial up
  - *Leased lines*: These are dedicated analog lines or digital lines. Dedicated digital lines are called digital data service (DDS) lines. A modem is used to connect to analog lines, and a Channel Service Unit/Data Service Unit or Digital Service Unit(CSU/DSU) is used to connect to digital lines. The DSU connects to the LAN and the CSU connects to the line.
  - *T Carrier lines*: Multiplexors are used to allow several channels on one line. The T1 line is basic T Carrier service. The available channels may be used separately for data or voice transmissions or they may be combined for more transmission bandwidth.

– T1 and T3 lines are the most common lines in use today. T1 and T2 lines can use standard copper wire. T3 and T4 lines require fibre-optic cable or other high-speed media. These lines may be leased partially called fractional T1 or fractional T3, which means a customer, can lease a certain number of channels on the line. A CSU/DSU and a bridge or router is required to connect to a T1 line.

– Integrated Services Digital Network (ISDN)-Comes in two types and converts analog signals to digital for transmission. It is a dial up service

a. Basic Rate ISDN (BRI)-Two 64Kbps B-channels with one 16Kbps D channel. The D-channel is used tor call control and setup.

b. Primary Rate ISDN (PRI)-23 B-channels and one D channel. A device resembling a modem (called an ISDN modem) is used to connect to ISDN. The computer and telephone line are plugged into it.

– *Switched-56*: A switched line similar to a leased line where customers pay for the time they use the line. Speed is 56Kbps. It is not dedicated and will not work to connect a WAN.

- *Asynchronous Transfer Mode (ATM)*: May be used over a variety of media with both baseband and broadband systems. It is used for audio, video, and data. It uses fixed length data packets of 53 8 bit bytes called cell switching. 5 bytes contain header information. The cell contains path information that the packet is to use. It uses hardware devices to perform the switching of the data. Speeds from 155Mbps to 622 Mbps are achieved. Error checking is done at the receiving device, not by ATM. A permanent virtual connection or circuit (PVC) is established. It may also use a switched virtual circuit (SVC). *Service classes*:

- Constant bit rate for data.
- Variable bit rate for audio or video.
- Connection less for data.
- Connection oriented for data.

ATM can be embedded in other protocols such as ATM-25, T1, T3, OC-1, OC-3, OC-12, and OC-48.

Some ATM technologies include:

- ATM-25-25Mbps speed.
- STS-3-155Mbps on fibre or category 5 cable.
- STS-12-620 Mbps on fibre cable for campus wide network.
- STS-48-2.2 Gbps on fibre cable on a MAN.
- STS-192-8.8 Gbps on fibre cable on intercity long distance. Phone companies normally use this.

Synchronous Optical Network (SONET)-A physical layer standard that defines voice, data, and video delivery methods over fibre optic media. It defines data rates in terms of optical carrier (OC) levels.

The transmission rate of OC-1 is 51.8 Mbps. Each level runs at a multiple of the first. The OC-5 data rate is 5 times 51.8 Mbps which is 259 Mbps.

SONET also defines synchronous transport signals (STS) for copper media which use the same speed scale of OC levels. STS-3 runs at the same speed of OC-3. Mesh or ring topology is used to support SONET. SONET uses multiplexing. The ITU has incorporated SONET into their Synchronous Digital Hierarchy (SDH) recommendations.

WAN Technology Comparisons

Terms

- *Circuit switching*: Physical switched connection.
- *Message switching*: A store and forward mechanism where messages are treated as individual units.
- *Packet switching*: Messages are broken down into smaller packets with individual destination information. Independent routing is used which allows the packets to use any route between the source send destination. Much RAM and processing power is required to support this switching type.
- *Data gram packet switching*: Uses independent paths.
- *Virtual circuit packet switching*: This is used for audio and video streaming. A set path is established between the source and destination and a connection-oriented service is made.

4

Network Services

FUTURE TRENDS

The evolution of communication networks from the current mix of analogue and digital access to full digital high capacity or broadband systems will occur. What is not clear is the exact path likely to be followed. Important networking issues will arise. With the availability at some sites of ATM technology, the use of FDDI for local networks at some institutions, and a large installed base of twisted pair Ethernet or Token Ring networks the migration path is not clear.

Access to ISDN, the availability of ISDN on PABX switches for internal use, the use of Metropolitan are networks to carry voice telephony and data traffic and the use of LANs for telephony all throw up important issues of network integration.

BROADBAND NETWORK SERVICES

B-ISDN (Broadband ISDN) with its ATM technology, going through the process of standardisation, is likely to be the most important candidate for the broadband multiservice network of the future. Much effort, worldwide, is being concentrated on building an understanding, through pilots projects, of this new technology. One of the main achievements of the ATM technology is its flexibility. It does not differentiate between various information characteristics, nor does it deal directly with end user time characteristics.

Users may not be satisfied with a low level direct access to the service offered by the ATM layer. To remedy this point an ATM (AAL) adaptation layer has been proposed. This layer will sit outside the pure ATM network and will provide the functionality not provided by the ATM network.

ATM Adaptation Layer Services

Currently there are four types of adaptation protocols progressing through the standardisation bodies: ITU-T, ETSI, and ATM forum. These are AAL-1, AAL-2, AAL-3/4 and AAL-5. Each AAL type is aimed at supporting specific communication requirements such as Connection Oriented, real-time issues, Variable Bit Rate etc.

The AAL-1 [ETS93c, ETS93e] offers a service that accepts Service Data Units (SDUs) at a fixed clock rate for transmission over the network and delivers them at the same clock rate in a Connection Oriented mode, also called isochronous service. The basic characteristic of the AAL-1 is the ability to offer a Constant Bit Rate (CBR) service. In addition it performs some basic error control mechanisms are performed which include sequence numbering in order to detect lost or mis-inserted ATM cells and Forward Error Correction (FEC). There are, currently, 3 services by AAL-1:

- The CBR circuit emulation service support the transmission of isochronous digital information. The CBR Circuit Emulation service provides two options: asynchronous circuit transport unstructured signals (2 Mbits/s or 34 Mbits/s); and synchronous circuit transport of 64kbits time slots.
- Video and Voice-band Signal Transport Services. The main differences between the CBR Circuit Emulation (CE) and these two services are the SDU size, one bit for CBR CE and one octet for the video and Voice-band Signal transport, and that for the Video Signal transport a specific FEC mechanism is foreseen. It is as yet undecided in how to handle voice signals in ATM cell stream with the problem of partly filled cells.

The AAL-2 protocol type should support Connection Oriented Variable Bit Rate (VBR) traffic and is intended to support transmission of VBR video codecs signal as an example. However, the specification process is at its early stages.

The AAL-3/4 protocol [ETS94b] offers a connection oriented service, not including any timing aspects, and is intended to support general data transmission applications. The basic functionality of the AAL-3/4 is SDU delimiting and multiplexing of higher layer SDUs on the ATM connection. From its incept the AAL-3/4 has been seen a provider of a Connectionless (CL) data service over the Connection Oriented ATM. Since it is an architectural requirement that the AAl should not deal with network layer issues, an additional layer is expected above the AAL-3/4 to offer a CL service. This CL protocol layer mainly provides the addressing functionalities in order to offer a Broadband Connectionless Data Service (BCDS) in ITU-T terminology or CBDS/SMDs in ETSI terminology. Therefore it makes no sense to consider the AAL type-3/4 service on its own.

The prime objective of the AAL-5 [ETS93d] is high speed transmission with reduced overhead. It is a compromise between overhead and functionality. The strong wish to develop a high speed AAL protocol type, which should support existing protocols, was the reason behind the specification of a reduced AAAL protocol type.

From a service point of view AAL-3/4 and AAL-5 offer the same layer functionality. The main differences between these two protocol types are: the AAL-5 performs minimum error control mechanisms in comparison to the AAL-3/4; they perform different mechanisms for SDU delimiting; and the AAL-5 does not offer a higher layer SDU multiplexing capability.

5

Network Devices System

It's important to remember that the firewall is only one entry point to your network. Modems, if you allow them to answer incoming calls, can provide an easy means for an attacker to sneak *around* (rather than *through*) your front door (or, firewall). Just as castles weren't built with moats only in the front, your network needs to be protected at all of its entry points.

SECURE MODEMS; DIAL-BACK SYSTEMS

If modem access is to be provided, this should be guarded carefully. The *terminal server*, or network device that provides dial-up access to your network needs to be actively administered, and its logs need to be examined for strange behaviour. Its passwords need to be strong — not ones that can be guessed. Accounts that aren't actively used should be disabled. In short, it's the easiest way to get into your network from remote: guard it carefully.

There are some remote access systems that have the feature of a two-part procedure to establish a connection. The first part is the remote user dialing into the system, and providing the correct userid and password. The system will then drop the connection, and call the authenticated user back at a known telephone number. Once the remote user's system answers that call, the connection is established, and the user is on the network. This works well for folks working at home, but can be problematic for users wishing to dial in from hotel rooms and such when on business trips.

Other possibilities include one-time password schemes, where the user enters his userid, and is presented with a "challenge," a string of between six and eight numbers. He types this challenge into a small device that he carries with him that looks like a calculator. He then presses enter, and a "response" is displayed on the LCD screen. The user types the response, and if all is correct, he login will proceed. These are useful devices for solving the problem of good passwords, without requiring dial-back access. However, these have their own problems, as they require the user to carry them, and they must be tracked, much like building and office keys.

No doubt many other schemes exist. Take a look at your options, and find out how what the vendors have to offer will help you *enforce your security policy effectively.*

CRYPTO-CAPABLE ROUTERS

A feature that is being built into some routers is the ability to use session encryption between specified routers. Because traffic travelling across the Internet can be seen by people in the middle who have the resources (and time) to snoop around, these are advantageous for providing connectivity between two sites, such that there can be secure routes.

VIRTUAL PRIVATE NETWORKS

Given the ubiquity of the Internet, and the considerable expense in private leased lines, many organizations have been building *VPNs* (Virtual Private Networks). Traditionally, for an organization to provide connectivity between a main office and a satellite one, an expensive data line had to be leased in order to provide direct connectivity between the two offices. Now, a solution that is often more economical is to provide both offices connectivity to the Internet. Then, using the Internet as the medium, the two offices can communicate.

The danger in doing this, of course, is that there is no privacy on this channel, and it's difficult to provide the other office access to "internal" resources without providing those resources to everyone on the Internet.

VPNs provide the ability for two offices to communicate with each other in such a way that it looks like they're directly connected over a private leased line. The session between them, although going over the Internet, is private (because the link is encrypted), and the link is convenient, because each can see each others' internal resources without showing them off to the entire world.

A number of firewall vendors are including the ability to build VPNs in their offerings, either directly with their base product, or as an add-on. If you have need to connect several offices together, this might very well be the best way to do it.

MODERN NETWORK DEVICES

Modern network devices are complex entities composed of both silicon and software. Thus, designing an efficient hardware platform is not, by itself, sufficient to achieve an effective, cost-efficient and operationally tenable product.

The control plane plays a critical role in the development of features and in ensuring device usability.

Although progress from the development of faster CPU boards and forwarding planes is visible, structural changes made in software are usually hidden, and while vendor collateral often offers a list of features in a carrier-class package, operational experiences may vary considerably. Products that have been through several generations of software releases provide the best examples of the difference made by the choice of OS. It is still not uncommon to find routers or switches that started life under older, monolithic software and later migrated to more contemporary designs. The positive effect on stability and operational efficiency is easy to notice and appreciate.

However, migration from one network operating system to another can pose challenges from non-overlapping feature sets, noncontiguous operational experiences and inconsistent software quality. These potential challenges make it is very desirable to build a control plane that can power the hardware products and features supported in both current and future markets. Developing a flexible, long-lasting and high-quality network OS provides a foundation that can gracefully evolve to support new needs in its height for up and down scaling, width for adoption across many platforms, and depth for rich integration of new features and functions. It takes time, significant investment and in-depth expertise.

Most of the engineers writing the early releases of Junos OS came from other companies where they had previously built network software. They had firsthand knowledge of what worked well, and what could be improved. These engineers found new ways to solve the limitations that they'd experienced in building the older operating systems.

Resulting innovations in Junos OS are significant and rooted in its earliest design stages. Still, to ensure that our products anticipate and fulfil the next generation of market requirements, Junos OS is periodically reevaluated to determine whether any changes are needed to ensure that it continues to provide the reliability, performance and resilience for which it is known.

Contemporary network operating systems are mostly advanced and specialized branches of POSIX-compliant software platforms and are rarely developed from scratch. The main reason for this situation is the high cost of developing a world-class operating system all the way from concept to finished product.

By adopting a general purpose OS architecture, network vendors can focus on routing-specific code, decrease time to market, and benefit from years of technology and research that went into the design of the original (donor) products.

FIRST-GENERATION OS: MONOLITHIC ARCHITECTURE

Typically, first-generation network operating systems for routers and switches were proprietary images running in a flat memory space, often directly from flash memory or ROM. While supporting multiple processes for protocols, packet handling and management, they operated using a cooperative, multitasking model in which each process would run to completion or until it voluntarily relinquished the CPU.

All first-generation network operating systems shared one trait: They eliminated the risks of running full-size commercial operating systems on embedded hardware. Memory management, protection and context switching were either rudimentary or nonexistent, with the primary goals being a small footprint and speed of operation.

Nevertheless, first-generation network operating systems made networking commercially viable and were deployed on a wide range of products. The downside was that these systems were plagued with a host of problems associated with resource management and fault isolation; a single runaway process could easily consume the processor or cause the entire system to fail. Such failures were not uncommon in the data networks controlled by older software and could be triggered by software errors, rogue traffic and operator errors.

Legacy platforms of the first generation are still seen in networks worldwide, although they are gradually being pushed into the lowest end of the telecom product lines.

SECOND-GENERATION OS: CONTROL PLANE MODULARITY

The mid-1990s were marked by a significant increase in the use of data networks worldwide, which quickly challenged the capacity of existing networks and routers. By this time, it had become evident that embedded platforms could run full-size commercial operating systems, at least on high-end hardware, but with one catch: They could not sustain packet forwarding with satisfactory data rates. A breakthrough solution was needed. It came in the concept of a hard separation between the control and forwarding plane—an approach that became widely accepted after the success of the industry's first application-specific integrated circuit (ASIC)-driven routing platform, the Juniper Networks M40. Forwarding packets entirely in silicon was proven to be viable, clearing the path for next generation network operating systems, led by Juniper with its Junos OS.

Today, the original M40 routers are mostly retired, but their legacy lives in many similar designs, and their blueprints are widely recognized in the industry as the second-generation reference architecture.

Second-generation network operating systems are free from packet switching and thus are focused on control plane functions. Unlike its first-generation counterparts, a second-generation OS can fully use the potential of multitasking, multithreading, memory management and context manipulation, all making systemwide failures less common. Most core and edge routers installed in the past few years are running second-generation operating systems, and it is these systems that are currently responsible for moving the bulk of traffic on the Internet and in corporate networks. However, the lack of a software data plane in second-generation operating systems prevents them from powering low-end devices without a separate (hardware) forwarding plane. Also, some customers cannot migrate from their older software easily because of compatibility issues and legacy features still in use.

These restrictions led to the rise of transitional (generation 1.5) OS designs, in which a first-generation monolithic image would run as a process on top of the second-generation scheduler and kernel, thus bridging legacy features with newer software concepts. The idea behind “generation 1.5” was to introduce some headroom and gradually move the functionality into the new code, while retaining feature parity with the original code base. Although interesting engineering exercises, such designs were not as feature-rich as their predecessors, nor as effective as their successors, making them of questionable value in the long term.

THIRD-GENERATION OS: FLEXIBILITY, SCALABILITY AND CONTINUOUS OPERATION

Although second-generation designs were very successful, the past 10 years have brought new challenges.

Increased competition led to the need to lower operating expenses and a coherent case for network software flexible enough to be redeployed in network devices across the larger part of the end-to-end packet path. From multiple terabit routers to Layer 2 switches and security appliances, the “best-in-class” catchphrase can no longer justify a splintered operational experience—true “network” operating systems are clearly needed. Such systems must also achieve continuous operation, so that software failures in the routing code, as well as system upgrades, do not affect the state of the network. Meeting this challenge requires availability and convergence characteristics that go far beyond the hardware redundancy available in second-generation routers.

Another key goal of third-generation operating systems is the capability to run with zero downtime (planned and unplanned). Drawing on the lesson learned from previous designs regarding the difficulty of moving from one OS to another, third-generation operating systems also should make the migration path completely transparent to customers. They must offer an evolutionary, rather than revolutionary upgrade experience typical to the retirement process of legacy software designs.

BASIC OS DESIGN CONSIDERATIONS

Choosing the right foundation (prototype) for an operating system is very important, as it has significant implications for the overall software design process and final product quality and serviceability. This importance is why OEM vendors sometimes migrate from one prototype platform to another midway through the development process, seeking a better fit. Generally, the most common transitions are from a proprietary to a commercial code base and from a commercial code base to an open-source software foundation.

Regardless of the initial choice, as networking vendors develop their own code, they get further and further away from the original port, not only in protocol-specific applications but also in the system area. Extensions such as control plane redundancy, in-service software upgrades and multi chassis operation require significant changes on all levels of the original design. However, it is highly desirable to continue borrowing content from the donor OS in areas that are not normally the primary focus of networking vendors, such as improvements in memory management, scheduling, multi core and symmetric multiprocessing (SMP) support, and host hardware drivers. With proper engineering discipline in place, the more active and peer-reviewed the donor OS is, the more quickly related network products can benefit from new code and technology.

This relationship generally explains another market trend—only two out of five network operating systems that emerged in the routing markets over the past 10 years used a commercial OS as a foundation.

Juniper's main operating system, Junos OS, is an excellent illustration of this industry trend. The basis of the Junos OS kernel comes from the FreeBSD UNIX OS, an open-source software system. The Junos OS kernel and infrastructure have since been heavily modified to accommodate advanced and unique features such as state replication, nonstop active routing and in-service software upgrades, all of which do not exist in the donor operating system. Nevertheless, the Junos OS tree can still be synchronized with the FreeBSD repository to pick the latest in system code, device drivers and development tool chains, which allows Juniper Networks engineers to concentrate on network-specific development.

Commercial Versus Open-Source Donor OS

The advantage of a more active and popular donor OS is not limited to just minor improvements—the cutting edge of technology creates new dimensions of product flexibility and usability.

Not being locked into a single-vendor framework and roadmap enables greater control of product evolution as well as the potential to gain from progress made by independent developers.

This benefit is evident in Junos OS, which became a first commercial product to offer hard resource separation of the control plane and a real-time software data plane. Juniper-specific extension of the original BSD system architecture relies on multicore CPUs and makes Junos OS the only operating system that powers both low-end software-only systems and high-end multiple-terabit hardware platforms with images built from the same code tree.

This technology and experience could not be created without support from the entire Internet-driven community. The powerful collaboration between leading individuals, universities and commercial organizations helps Junos OS stay on the very edge of operating system development. Further, this collaboration works both ways:

Juniper donates to the free software movement, one example being the Juniper Networks FreeBSD/MIPS port.

Functional Separation and Process Scheduling

Multiprocessing, functional separation and scheduling are fundamental for almost any software design, including network software. Because CPU and memory are shared resources, all running threads and processes have to access them in a serial and controlled fashion. Many design choices are available to achieve this goal, but the two most important are the memory model and the scheduling discipline.

Memory Model

The memory model defines whether processes (threads) run in a common memory space. If they do, the overhead for switching the threads is minimal, and the code in different threads can share data via direct memory pointers.

The downside is that a runaway process can cause damage in memory that does not belong to it.

In a more complex memory model, threads can run in their own virtual machines, and the operating system switches the context every time the next thread needs to run. Because of this context switching, direct communication between threads is no longer possible and requires special Inter Process Communication (IPC) structures such as pipes, files and shared memory pools.

Scheduling Discipline

Scheduling choices are primarily between cooperative and preemptive models, which define whether thread switching happens voluntarily. A cooperative multitasking model allows the thread to run to completion, and a preemptive design ensures that every thread gets access to the CPU regardless of the state of other threads.

Virtual Memory/Preemptive Scheduling Programming Model

Virtual memory with preemptive scheduling is a great design choice for properly constructed functional blocks, where interaction between different modules is limited and well defined. This technique is one of the main benefits of the second-generation OS designs and underpins the stability and robustness of contemporary network operating systems. However, it has its own drawbacks.

Notwithstanding the overhead associated with context switching, consider the interaction between two threads, A and B, both relying on the common resource R. Because threads do not detect their relative scheduling in the preemptive model, they can actually access R in a different order and with varying intensity. For example, R can be accessed by A, then B, then A, then A and then B again. If thread B modifies resource R, thread A may get different results at different times—and without any predictability. For instance, if R is an interior gateway protocol (IGP) next hop, B is an IGP process, and A is a BGP process, then BGP route installation may fail because the underlying next hop was modified midway through routing table modification. This scenario would never happen in the cooperative multitasking model, because the IGP process would release the CPU only after it finishes the next-hop maintenance. This problem is well researched and understood within software design theory, and special solutions such as resource locks and synchronization primitives are easily available in nearly every operating system. However, the effectiveness of IPC depends greatly on the number of interactions between different processes. As the number of interacting processes increases, so does the number of IPC operations. In a carefully designed system, the number of IPC operations is proportional to the number of processes (N). In a system with extensive IPC activity, this number can be proportional to N2.

Exponential growth of an IPC map is a negative trend not only because of the associated overhead, but because of the increasing number of unexpected process interactions that may escape the attention of software engineers.

In practice, overgrown IPC maps result in systemwide "IPC meltdowns" when major events trigger intensive interactions. For instance, pulling a line card would normally affect interface management, IGP, exterior gateway protocol and traffic engineering processes, among others. When interprocess interactions are not well contained, this event may result in locks and tight loops, with multiple threads waiting on each other and vital system operations such as routing table maintenance and IGP computations temporarily suspended. Such defects are signatures of improper modularization, where similar or heavily interacting functional parts do not run as one process or one thread.

The right question to ask is, "Can a system be too modular?" The conventional wisdom says, "Yes." Excessive modularity can bring long-term problems, with code complexity, mutual locks and unnecessary process interdependencies. Although none of these may be severe enough to halt development, feature velocity and scaling parameters can be affected. Complex process interactions make programming for such a network OS an increasingly difficult task.

On the other hand, the cooperative multitasking, shared memory paradigm becomes clearly suboptimal if unrelated processes are influencing each other via the shared memory pool and collective restartability. A classic problem of first-generation operating systems was systemwide failure due to a minor bug in a nonvital process such as SNMP or network statistics. Should such an error occur in a protected and independently restartable section of system code, the defect could easily be contained within its respective code section.

This brings us to an important conclusion. No fixed principle in software design fits all possible situations. Ideally, code design should follow the most efficient paradigm and apply different strategies in different parts of the network OS to achieve the best marriage of architecture and function.

This approach is evident in Junos OS, where functional separation is maintained so that cooperative multitasking and preemptive scheduling can both be used effectively, depending on the degree of IPC containment between functional modules.

Generic Kernel Design

Kernels normally do not provide any immediately perceived or revenue-generating functionality. Instead, they perform housekeeping activities such as memory allocation and hardware management and other system-level tasks. Kernel threads are likely the most often run tasks in the entire system. Consequently, they have to be robust and run with minimal impact on other processes.

In the past, kernel architecture largely defined the operating structure of the entire system with respect to memory management and process scheduling. Hence, kernels were considered important differentiators among competing designs.

Historically, the disputes between the proponents and opponents of lightweight versus complex kernel architectures came to a practical end when most operating systems became functionally decoupled from their respective kernels.

Once software distributions became available with alternate kernel configurations, researchers and commercial developers were free to experiment with different designs.

For example, the original Carnegie-Mellon Mach microkernel was originally intended to be a drop-in replacement for the kernel in BSD UNIX and was later used in various operating systems, including mkLinux and GNU FSF projects. Similarly, some software projects that started life as purely microkernel-based systems later adopted portions of monolithic designs.

Over time, the radical approach of having a small kernel and moving system functions into the user-space processes did not prevail. A key reason for this was the overhead associated with extra context switches between frequently executed system tasks running in separate memory spaces.

Furthermore, the benefits associated with restart ability of essentially all system processes proved to be of limited value, especially in embedded systems. With the system code being very well tested and limited to scheduling, memory management and a handful of device drivers, the potential errors in kernel subsystems are more likely to be related to hardware failures than to software bugs. This means, for example, that simply restarting a faulty disk driver is unlikely to help the routing engine stay up and running, as the problem with storage is likely related to a hardware failure (for example, uncorrectable fault in a mass storage device or system memory bank).

Another interesting point is that although both monolithic and lightweight kernels were widely studied by almost all operating system vendors, few have settled on purist implementations.

For example, Apple's Mac OS X was originally based on microkernel architecture, but now runs system processes, drivers and the operating environment in BSD-like subsystems. Microsoft NT and derivative operating systems also went through multiple changes, moving critical performance components such as graphical and I/O subsystems in and out of the system kernel to find the right balance of stability, performance and predictability. These changes make NT a hybrid operating system. On the other hand, freeware development communities such as FSF, FreeBSD and NetBSD have mostly adopted monolithic designs (for example, Linux kernel) and have gradually introduced modularity into selected kernel sections (for example, device drivers). So what difference does kernel architecture make to routing and control?

RECEIVING MULTICAST MESSAGES

Receiving a multicast message is a bit tricker. Since multicast messages are generally ignored unless someone has explicitly registered an interest in a particular address, there is a bit of setup that needs to be done.

Here's the code for receiving multicast messages:

```
require 'socket'
require 'ipaddr'
```

```
MULTICAST_ADDR = "225.4.5.6"
PORT = 5000
ip = IPAddr.new(MULTICAST_ADDR).hton + IPAddr.new("0.0.0.0").hton
sock = UDPSocket.new
sock.setsockopt(Socket::IPPROTO_IP, Socket::IP_ADD_MEMBERSHIP, ip)
sock.bind(Socket::INADDR_ANY, PORT)
loop do
    msg, info = sock.recvfrom(1024)
    puts "MSG: #{msg} from #{info[2]} (#{info[3]})/#{info[1]} len #{msg.size}"
end
```

The tricky part was figuring out the right setsockopt options and values needed to register interest in our multicast address. I had to do a little reading in the Unix man pages on the C level setsockopt() function call.

The third option to the C function is a structure that contains two 4-byte IP addresses. The first IP address is the multicast address, and the second IP address is the address of the local host adapter that we wish to use to listen for the multicast. The 0.0.0.0 address means use any of the local network adapters.

IPAddr handles parsing the human readable form of the IP address and returns a string of 4 bytes in the order needed by the C level setsockopt() function.

USAGE

Save the above code in files named send.rb and rcv.rb. In one console window, type:

RUBY RCV.RB

In another console window on the same or different machine (on the same local network), type:

RUBY SEND.RB THIS IS A TEST.

UPDATE

I added the Time To Live option on send.

POINT-TO-POINT COMMUNICATIONS

A point-to-point connection is a dedicated communication link between two systems or processes. Think of a wire that directly connects two systems. The systems use that wire exclusively to communicate. The opposite of point-to-point communications is broadcasting, where one system transmits to many. A telephone call is a circuit-oriented, point-to-point link between two phones. However, calls are usually multiplexed across telephone company trunks; so, while the circuit itself may be virtual, the users are engaging in a point-to-point communication session.

An end-to-end connection refers to a connection between two systems across a switched network. For example, the Internet is made up of a mesh of routers. Packets follow a hop-by-hop path from one router to the next to reach their destinations.

Each hop consists of a physical point-to-point link between routers. Therefore, a routed path consists of multiple point-to-point links. In the ATM and frame relay environment, the end-to-end path is called a virtual circuit that crosses a predefined set of point-to-point links.

A shared LAN such as Ethernet provides a form of point-to-point communications. Keep in mind that on shared LANs, all nodes listen to signals on the cable, so broadcasting is supported. However, when one node addresses frames to another node and only that node receives the frames, one could say that the two nodes are engaged in point-to-point communications across a shared medium.

Point-to-multipoint connections are possible over multidrop links. A mainframe and its terminals is an example. The device that provides the multipoint connection is usually an intelligent controller that manages the flow of information from the multiple devices attached to it.

Point-to-point communications is defined in the physical and data link layers of the OSI protocol stack.

Process one (myrank = 1) receives this message with the receive operation MPI_RECV. The message to be received is selected according to the value of its envelope, and the message data is stored into the receive buffer. In the example above, the receive buffer consists of the storage containing the string message in the memory of process one. The first three parameters of the receive operation specify the location, size and type of the receive buffer. The next three parameters are used for selecting the incoming message. The last parameter is used to return information on the message just received.

The next sections describe the blocking send and receive operations. We discuss send, receive, blocking communication semantics, type matching requirements, type conversion in heterogeneous environments, and more general communication modes. Non-blocking communication is addressed next, followed by channel-like constructs and send-receive operations. We then consider general datatypes that allow one to transfer efficiently heterogeneous and non-contiguous data. We conclude with the description of calls for explicit packing and unpacking of messages. The basic communication mechanism of MPI is the transmittal of data between a pair of processes, one side sending, the other, receiving. We call this ``point to point communication.'' Almost all the constructs of MPI are built around the point to point operations and so this chapter is fundamental. It is also quite a long chapter since: there are many variants to the point to point operations; there is much to say in terms of the semantics of the operations; and related topics, such as probing for messages, are explained here because they are used in conjunction with the point to point operations.

MPI provides a set of send and receive functions that allow the communication of typedtyped data data with an associated tag.tagmessage tag Typing of the message contents is necessary for heterogeneous support - the type information is needed so that correct data representation conversions can be performed as data is sent from one architecture to another. The tag allows selectivity of messages at the receiving end: one can receive on a particular tag, or one can wild-card this quantity, allowing reception of messages with any tag. Message selectivity on the source process of the message is also provided.

A fragment of C code appears in Example for the example of process 0 sending a message to process 1. The code executes on both process 0 and process 1. Process 0 sends a character string using MPI_Send(). The first three parameters of the send call specify the data to be sent: the outgoing data is to be taken frommsg; it consists of strlen(msg)+1 entries, each of type MPI_CHAR (The string "Hello there" contains strlen(msg)=11 significant characters. In addition, we are also sending the tex2html_html_special_mark_quot''" string terminator character). The fourth parameter specifies the message destination, which is process 1. The fifth parameter specifies the message tag. Finally, the last parameter is a communicatorcommunicator that specifies a communication domaincommunication domain for this communication. Among other things, a communicator serves to define a set of processes that can be contacted. Each such process is labeled by a process rank.rank Process ranks are integers and are discovered by enquiry to a communicator (see the call to MPI_Comm_rank()). MPI_COMM_ WORLDMPI_ COMM_WORLD is a default communicator provided upon start-up that defines an initial communication domain for all the processes that participate in the computation. Much more will be said about communicators in Chapter .

We have already said rather a lot about a simple transmittal of data from one process to another, but there is even more. To understand why, we examine two aspects of the communication: the semantics semantics of the communication primitives, and the underlying protocols that protocols implement them. Consider the previous example, on process 0, after the blocking send has completed. The question arises: if the send has completed, does this tell us anything about the receiving process? Can we know that the receive has finished, or even, that it has begun?

Such questions of semantics are related to the nature of the underlying protocol implementing the operations. If one wishes to implement a protocol minimizing the copying and buffering of data, the most natural semantics might be the ``rendezvous''rendezvous version, where completion of the send implies the receive has been initiated (at least). On the other hand, a protocol that attempts to block processes for the minimal amount of time will necessarily end up doing more buffering and copying of data and will have ``buffering'' semantics.buffering The trouble is, one choice of semantics is not best for all applications, nor is it best for all architectures. Because the primary goal of MPI is to standardize the operations, yet not sacrifice performance, the decision was made to include all the major choices for point to point semantics in the standard. The above complexities are manifested in MPI by the existence of modesmodes for point to point communication. Both blocking and non-blocking communications have modes. The mode allows one to choose the semantics of the send operation and, in effect, to influence the underlying protocol of the transfer of data.

In standard modestandard mode the completion of the send does not necessarily mean that the matching receive has started, and no assumption should be made in the application programme about whether the out-going data is buffered by MPI. In buffered mode buffered mode the user can guarantee that a certain amount of buffering space is available. The catch is that the space must be explicitly provided by the application programme. In synchronous mode synchronous mode a rendezvous semantics between sender and receiver is used. Finally, there is ready mode. ready mode This allows the user to exploit extra knowledge to simplify the protocol and potentially achieve higher performance. In a ready-mode send, the user asserts that the matching receive already has been posted. Modes are covered in Section .

INTERNET PROTOCOL ADDRESSING

An IP (Internet Protocol) address is a unique identifier for a node or host connection on an IP network. An IP address is a 32 bit binary number usually represented as 4 decimal values, each representing 8 bits, in the range 0 to 255 (known as octets) separated by decimal points. This is known as "dotted decimal" notation. Every IP address consists of two parts, one identifying the network and one identifying the node. The Class of the address and the subnet mask determine which part belongs to the network address and which part belongs to the node address.

There are 5 different address classes. You can determine which class any IP address is in by examining the first 4 bits of the IP address.

- Class A addresses begin with 0xxx, or 1 to 126 decimal.
- Class B addresses begin with 10xx, or 128 to 191 decimal.
- Class C addresses begin with 110x, or 192 to 223 decimal.
- Class D addresses begin with 1110, or 224 to 239 decimal.
- Class E addresses begin with 1111, or 240 to 254 decimal.
- Addresses beginning with 01111111, or 127 decimal, are reserved for loopback and for internal testing on

a local machine; Class D addresses are reserved for multicasting; Class E addresses are reserved for future use. They should not be used for host addresses.

PRIVATE SUBNETS

There are three IP network addresses reserved for private networks. The addresses are 10.0.0.0, Subnet Mask 255.0.0.0, 172.16.0.0, Subnet Mask 255.240.0.0, and 192.168.0.0, Subnet Mask 255.255.0.0. These addresses are also notated 10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16; this notation will be explained later in this tutorial. They can be used by anyone setting up internal IP networks, such as a lab or home LAN behind a NAT or proxy server or a router. It is always safe to use these because routers on the Internet by default will never forward packets coming from these addresses. These addresses are defined in RFC 1918.

IP ADDRESSING AND SUBNETTING

This document gives you basic information needed to configure your router for routing IP, such as how addresses are broken down and how subnetting works. You learn how to assign each interface on the router an IP address with a unique subnet. There are many examples to help tie everything together.

Prerequisites

Requirements

There are no specific prerequisites for this document.

Additional Information

If definitions are helpful to you, use these vocabulary terms to get you started:

- Address-The unique number ID assigned to one host or interface in a network.
- Subnet-A portion of a network sharing a particular subnet address.
- Subnet mask-A 32-bit combination used to describe which portion of an address refers to the subnet and which part refers to the host.
- Interface-A network connection.

If you have already received your legitimate address(es) from the Internet Network Information Center (InterNIC), you are ready to begin. If you do not plan to connect to the Internet, Cisco strongly suggests that you use reserved addresses from RFC 1918 .

Network Masks

A network mask helps you know which portion of the address identifies the network and which portion of the address identifies the node. Class A, B, and C networks have default masks, also known as natural masks, as shown here:

Class A: 255.0.0.0

Class B: 255.255.0.0

Class C: 255.255.255.0

An IP address on a Class A network that has not been subnetted would have an address/mask pair similar to:

8.20.15.1 255.0.0.0. To see how the mask helps you identify the network and node parts of the address, convert the address and mask to binary numbers.

8.20.15.1 = 00001000.00010100.00001111.00000001

255.0.0.0 = 11111111.00000000.00000000.00000000

Once you have the address and the mask represented in binary, then identifying the network and host ID is easier. Any address bits which have corresponding mask bits set to 1 represent the network ID. Any address bits that have corresponding mask bits set to 0 represent the node ID.

8.20.15.1 = 00001000.00010100.00001111.00000001

255.0.0.0 = 11111111.00000000.00000000.00000000

net id | host id

netid = 00001000 = 8

hostid = 00010100.00001111.00000001 = 20.15.1

6

Networking Hardware

Networking hardware includes all computers, peripherals, interface cards and other equipment needed to perform data-processing and communications within the network.

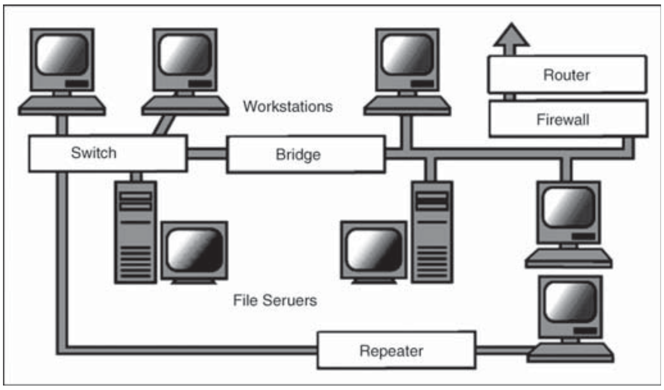

FILE/NETWORK SERVERS

One or more network servers is a part of nearly every local area network. These are very fast computers with a large amount of RAM and storage space, along with a one or more fast network interface card(s). The network operating system provides tools to share server resources and information with network users. A sophisticated permissions-handling system is included, so that access to sensitive information can be carefully tailored to the needs of the users. For small networks, a singe network server may provide access control, file sharing, printer sharing, e-mail, database, and other services. The network server may be responding to requests from many network users simultaneously. For example, it may be asked to load a word processor program to one workstation, receive a database file from another workstation, and store an e-mail message during the same time period. This requires a computer that can store and quickly share large amounts of information. When configuring such a server, budget is usually the controlling factor.

The following guidelines should be followed:

- Fastest processor(s)
- Large amount of RAM
- Multiple large, fast hard drives
- Extra expansion slots
- Fast network interface card(s).

Optionally (if no other such devices are available on the network):

- A RAID (Redundant Array of Inexpensive Disks) to preserve large amounts of data (even after a disk failure)

- A back-up unit (*i.e.* DAT tape drive, removable hard drives, or CD/DVD/BluRay burner).

WORKSTATIONS

Computers that humans use are broadly categorized as workstations. A typical workstation is a computer that is configured with a network interface card, networking software, and the appropriate cables. Workstations do not necessarily need large storage hard drives, because files can be saved on the file server. Almost any computer can serve as a network workstation.

LAPTOPS/MOBILE DEVICES

Laptops and other mobile devices are becoming more and more common. These devices typically have modest internal storage, but enough power to serve as a workstation for users on the go. These machines nearly always have a wireless adapter to allow quick network connections without cumbersome cabling. In a school environment with good wireless coverage, a mobile device user can move about the campus freely, and remain continuously connected to the network.

NETWORK INTERFACE CARDS

The network interface card (NIC) provides the physical connection between the network and the computer workstation. Most NICs are internal, and they are included in the purchase of most computers. Network interface cards are a major factor in determining the speed and performance of a network. It is a good idea to use the fastest network card available for the type of workstation you are using. The most common network interface connections are Ethernet cards and wireless adapters.

ETHERNET CARDS

Ethernet cards are usually included with a computer, although additional ethernet cards can be purchased and installed on most computers,. Ethernet cards can contain connections for either coaxial or twisted pair cables. If it is designed for coaxial cable, the connection will be BNC. If it is designed for twisted pair, it will have a RJ-45 connection. Some Ethernet cards also contain an AUI connector. This can be used to attach coaxial, twisted pair, or fibre optics cable to an Ethernet card. When this method is used there is always an external transceiver attached to the workstation. Only the RJ-45 connector is found on most modern ethernet cards.

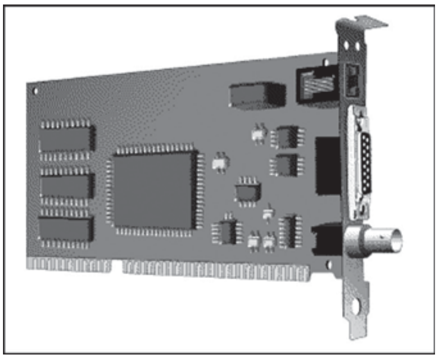

Figure : Ethernet Card.

WIRELESS ADAPTERS

Wireless adapters are found in most portable devices, such as laptops, smart phones, and tablet devices. External wireless adapters can be purchased and installed on most computers having an open USB (Universal Serial Bus) port, or unused expansion slot.

SWITCHES

An ethernet switch is a device that provides a central connection point for cables from workstations, servers, and peripherals. In a star topology, twisted-pair wire is run from each workstation to a central switch/hub. Most switches are active, that is they electrically amplify the signal as it moves from one device to another. The predecessor of the switch was the hub, which broadcasted all inbound packets out all ports of the device, creating huge amounts of unnecessary network traffic. Modern switches build a port map of all IP address which respond on each port, and only broadcasts on all ports when it doesn't have a packet's target IP address already in its port map. Switches are:

- Usually configured with 8, 12, or 24 RJ-45 ports
- Often used in a star or tree topology
- Available as "managed" or "unmanaged", with the later less expensive, but adequate for smaller networks
- Direct replacements for hubs, immediately reducing network traffic in most networks

- Usually installed in a standardized metal rack that also may store network servers, bridges, or routers.

REPEATERS

Since a signal loses strength as it passes along a cable, it is often necessary to boost the signal with a device called a repeater. The repeater electrically amplifies the signal it receives and rebroadcasts it. Repeaters can be separate devices or they can be incorporated into a concentrator. They are used when the total length of your network cable exceeds the standards set for the type of cable being used. A good example of the use of repeaters would be in a local area network using a star topology with unshielded twisted-pair cabling. The length limit for unshielded twisted-pair cable is 100 meters. The most common configuration is for each workstation to be connected by twisted-pair cable to a multi-port active concentrator. The concentrator amplifies all the signals that pass through it allowing for the total length of cable on the network to exceed the 100 meter limit.

BRIDGES

A bridge is a device that allows you to segment a large network into two smaller, more efficient networks. If you are adding to an older wiring scheme and want the new network to be up-to-date, a bridge can connect the two.

A bridge monitors the information traffic on both sides of the network so that it can pass packets of information to the correct location. Most bridges can "listen" to the network and automatically figure out the address of each computer on both sides of the bridge. The bridge can inspect each message and, if necessary, broadcast it on the other side of the network.

The bridge manages the traffic to maintain optimum performance on both sides of the network. You might say that the bridge is like a traffic cop at a busy intersection during rush hour. It keeps information flowing on both sides of the network, but it does not allow unnecessary traffic through. Bridges can be used to connect different types of cabling, or physical topologies. They must, however, be used between networks with the same protocol.

ROUTERS

Routers are the traffic directors of the global internet. All routers maintain complex routing tables which allow them to determine appropriate paths for packets destined for any address. Routers communicate with each other, and forward network packets out of or into a network. Here's an example:

You want to search for something on the internet using a search engine. You open a browser on your workstation. The browser opens to a blank page (not usually the default, but appropriate for this example). You type "http://

www.google.com" into the URL (Universal Resource Locator) address line of the browser. The browser software packages up the URL you typed, and sends it with a request for an IP address to the DNS (Domain Name Server) that has been set in your network adapter's configuration. The domain server returns an IP, such as 74.125.67.103 (actual address returned by DNS for google.com on June 7th, 2011). The browser ships the request for that IP address off to the network card, which bundles the request into an ethernet packet, destined for 74.125.67.103. The network card sends the packet to the gateway of your network, which opens the header of the packet, and makes a determination that the packet is travelling out of your network, in search of 74.125.67.103. Your network's router has routing tables which it has been building from communicating with other routers, and potentially augmented with "static routes", which are specific paths added by your network's administrators to make the task of accessing certain networks easier, or faster, or in some cases, not possible. In this case, I find that my router knows about another router at my ISP (Internet Service Provider), which in turn has several more routers that are all on networks of which I am just a small node, much like finding an atom of a molecule of a piece of dust on a rock on a moon of a planet of a sun of a galaxy of the universe. In any case, the packet gets passed from router to router, each time moving out of the subnets of the packet sender, towards a router that will know where the desired server is. The packet finally reaches the router of the network at 74.125.67.103, which dutifully delivers the packet to the server at that IP address. The server carefully crafts a response, and sends a reply back, which follows the same process to get the response “Yes. Go ahead” back to the requester. Whew. And that’s just the initial request.

While bridges know the addresses of all computers on each side of the network, routers know the addresses other routers which in turn know about their own networks. Routers can even “listen” to entire networks to determine which sections are busiest — they can then redirect data around those sections until traffic congestion clears.

So, routers are network gateways. They move network packets from one network to another, and many can convert from one network protocol to another as necessary. Routers select the best path to route a message, based on the destination address of the packet. The router can direct traffic to prevent head-on collisions, and is smart enough to know when to direct traffic along back roads and shortcuts.

If you have a school LAN that you want to connect to the Internet, you will need to purchase a router. In this case, the router serves as the forwarder between the information on your LAN and the Internet. It also determines the best route to send the data over the Internet.

FIREWALLS

A firewall is a networking device that is installed at the entrance to a LAN when connecting a network together, particularly when connecting a private network to a public network, such as the internet. The firewall uses rules to filter traffic into and out of the private network, to protect the private network users and data from malevolent hackers.

Firewalls are either hardware or software, depending on their intended use. A firewall used to protect a network is a hardware device that should be installed in the network between the router and the network. Almost all hardware firewalls will have at least two ports, labelled "Trusted" and "Untrusted". These terms imply the true nature of the firewall's responsibility to the private network. The public network is connected to the untrusted network port, and the private network is connected to the trusted port.

Firewall rules are usually simple, consisting of a verb, either allow or deny, the direction of the traffic, either inbound or outbound, and an address or other network traffic identifier. Firewall rules are cumulative, so general rules may be specified, and exceptions added as necessary. Some examples are:

- Allow outbound all (all private network users can do anything on the public network)
- Deny inbound all (default setting to prevent all traffic from the public or untrusted port, to the private port)

- Allow inbound port 80 (allow internet web traffic to come into network to find web servers)
- Allow inbound port 80 destined to 170.200.201.25 (allow inbound web traffic to a specific web server on your private network)
- Deny inbound from 201.202.1.1/24 (deny all inbound traffic from a specific IP address or range of addresses).

Software firewalls are commonly included in modern workstation and server operating systems. They operate in a similar way as hardware firewalls, except that they filter traffic in and out of the machine itself. These software firewalls are typically unnoticed by machine users, and only need attention occasionally when an internet-connected application don't work as expected. The software firewall should always be considered a "suspect" in such cases. The problem is easily resolved, by setting an exception rule in the firewall for the software that is attempting to communicate.

HUBS AND SWITCHES

As you can see, computers use network cards to send and receive data. The data is transmitted over Ethernet cables. However, you normally can't just run an Ethernet cable between two PCs and call it a network.

In this day and age of high speed Internet access being almost universally available, you tend to hear the term broadband thrown around a lot. Broadband is a type of network in which data is sent and received across the same wire. In contrast, Ethernet uses Baseband communications. Baseband uses separate wires for sending and receiving data. What this means is that if one PC is sending data across a particular wire within the Ethernet cable, then the PC that is receiving the data needs to have the wire redirected to its receiving port.

You can actually network two PCs together in this way. You can create what is known as a cross over cable. A cross over cable is simply a network cable that has the sending and receiving wires reversed at one end, so that two PCs can be linked directly together.

The problem with using a cross over cable to build a network is that the network will be limited to using no more and no less than two PCs. Rather than using a cross over cable, most networks use normal Ethernet cables that do not have the sending and receiving wires reversed at one end.

Of course the sending and receiving wires have to be reversed at some point in order for communications to succeed. This is the job of a hub or a switch. Hubs are starting to become extinct, but I want to talk about them any way because it will make it easier to explain switches later on.

There are different types of hubs, but generally speaking a hub is nothing more than a box with a bunch of RJ-45 ports.

Each computer on a network would be connected to a hub via an Ethernet cable.

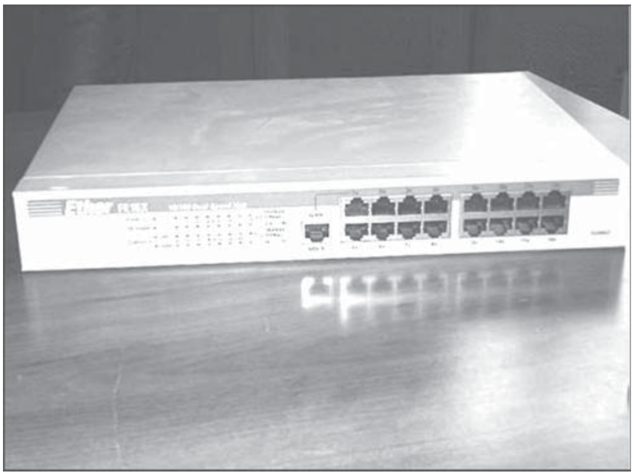

Figure : A Hub is a Device that Acts as a Central Connection Point for Computers on a Network.

A hub has two different jobs. Its first job is to provide a central point of connection for all of the computers on the network. Every computer plugs into the hub (multiple hubs can be daisy chained together if necessary in order to accommodate more computers). The hub's other job is to arrange the ports in such a way so that if a PC transmits data, the data is sent over the other computer's receive wires. Right now you might be wondering how data gets to the correct destination if more than two PCs are connected to a hub. The secret lies in the network card. Each Ethernet card is programmed at the factory with a unique Media Access Control (MAC) address. When a computer on an Ethernet network transmits data across an Ethernet network containing PCs connected to a hub, the data is actually sent to every computer on the network. As each computer receives the data, it compares the destination address to its own MAC address. If the addresses match then the computer knows that it is the intended recipient, otherwise it ignores the data.

As you can see, when computers are connected via a hub, every packet gets sent to every computer on the network. The problem is that any computer can send a transmission at any given time. Have you ever been on a conference call and accidentally started to talk at the same time as someone else? This is the same thing that happens on this type of network.

When a PC needs to transmit data, it checks to make sure that no other computers are sending data at the moment. If the line is clear, it transmits the necessary data. If another computer tries to communicate at the same time though, then the packets of data that are travelling across the wire collide and are destroyed (this is why this type of network is sometimes referred to as a collision domain). Both PCs then have to wait for a random amount of time and attempt to retransmit the packet that was destroyed.

As the number of PCs on a collision domain increases, so does the number of collisions. As the number of collisions increase, network efficiency is decreased. This is why switches have almost completely replaced hubs.

A switch, performs all of the same basic tasks as a hub. The difference is that when a PC on the network needs to communicate with another PC, the switch uses a set of internal logic circuits to establish a dedicated, logical path between the two PCs. What this means is that the two PCs are free to communicate with each other, without having to worry about collisions.

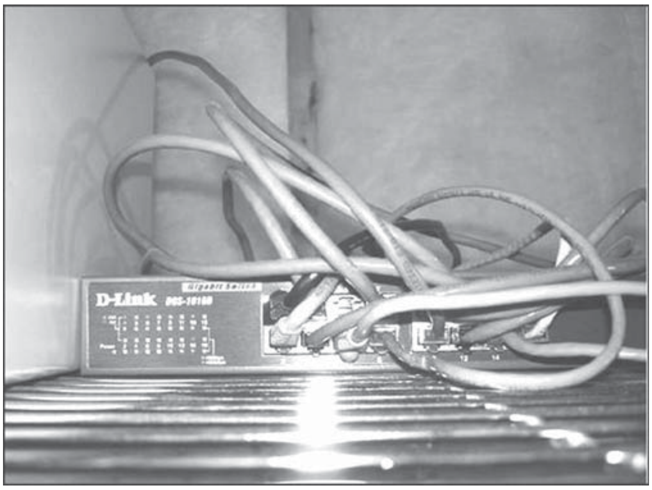

Figure : A Switch Looks a lot Like a Hub, but Performs very Differently.

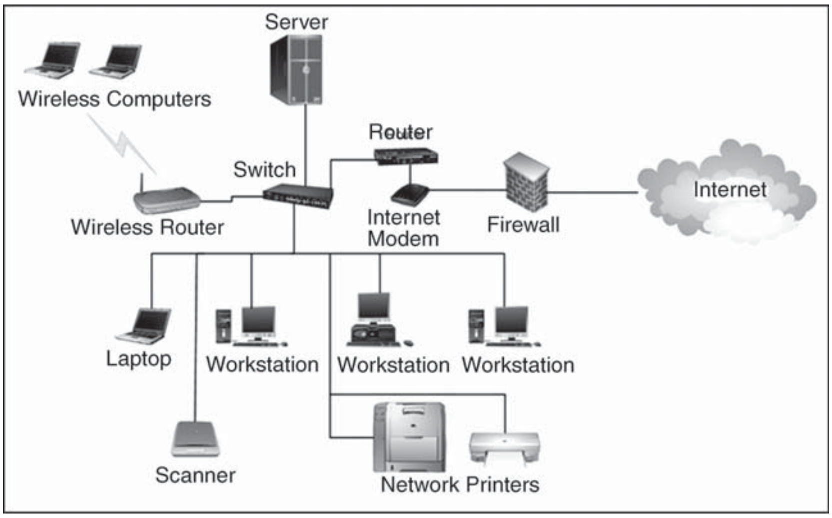

Switches greatly improve a network's efficiency. Yes, they eliminate collisions, but there is more to it than that. Because of the way that switches work, they can establish parallel communications paths.

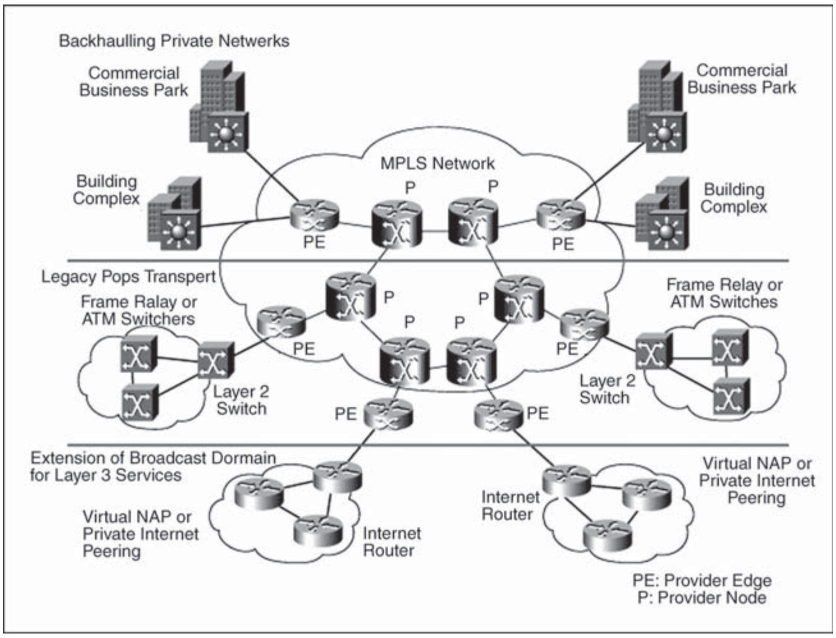

For example, just because computer A is communicating with computer B, there is no reason why computer C can't simultaneously communicate with computer D. In a collision domain, these types of parallel communications would be impossible because they would result in collisions.

7

Multimedia Messaging Service

Multimedia Messaging Service, or MMS, is a standard way to send messages that include multimedia content to and from mobile phones. It extends the core SMS (Short Message Service) capability that allowed exchange of text messages only up to 160 characters in length. The most popular use is to send photographs from camera-equipped handsets, although it is also popular as a method of delivering news and entertainment content including videos, pictures, text pages and ringtones. The standard is developed by the Open Mobile Alliance (OMA), although during development it was part of the 3GPP and WAP groups.

History

The immediate predecessor to the MMS is the Japanese picture messaging system Sha-Mail introduced by J-Phone in 2001. It validated the concept of camera phone users willing to send picture messages from one phone to another. Early MMS deployments were plagued by technical issues and frequent consumer disappointments, such as having sent an MMS message, receiving a confirmation it had been sent, being billed for the MMS message, to find that it had not been delivered to the intended recipient. Pictures would often arrive in the wrong formats, and other media elements might be removed such as a video clip arriving without its sound. At the MMS World Congress in 2004 in Vienna, all European mobile operator representatives who had launched MMS, admitted their MMS services were not making money for their networks. Also on all networks at the time, the most common uses were various adult oriented services that had been deployed using MMS. China was one of the early markets to make MMS a major commercial success partly as the penetration rate of personal computers was modest but MMS-capable cameraphones spread rapidly. The chairman and CEO of China Mobile said at the GSM Association Mobile Asia Congress in 2009 that MMS in China is now a mature service on par with SMS text messaging. Europe's most advanced MMS market has been Norway and in 2008 the Norwegian MMS usage level had passed 84% of all mobile phone subscribers. Norwegian mobile subscribers average one MMS sent per week. By 2008 worldwide MMS usage level had passed 1.3 billion active users who generated

50 billion MMS messages and produced annual revenues of 26 billion dollars.

Technical Description

MMS messages are delivered in a completely different way from SMS. The first step is for the sending device to encode the multimedia content in a fashion similar to sending a MIME e-mail (MIME content formats are defined in the MMS Message Encapsulation specification). The message is then forwarded to the carrier's MMS store and forward server, known as the MMSC. If the receiver is on another carrier, the relay forwards the message to the recipient's carrier using the Internet. Once the MMSC has received a message, it first determines whether the receiver's handset is "MMS capable", that is it supports the standards for receiving MMS. If so, the content is extracted and sent to a temporary storage server with an HTTP front-end. An SMS "control message" containing the URL of the content is then sent to the recipient's handset to trigger the receiver's WAP browser to open and receive the content from the embedded URL.

Several other messages are exchanged to indicate status of the delivery attempt. Before delivering content, some MMSCs also include a conversion service that will attempt to modify the multimedia content into a format suitable for the receiver. This is known as "content adaptation". If the receiver's handset is not MMS capable, the message is usually delivered to a web based service from where the content can be viewed from a normal internet browser. The URL for the content is usually sent to the receiver's phone in a normal text message. This behaviour is usually known as the "legacy experience" since content can still be received by a phone number, even if the phone itself does not support MMS. The method for determining whether a handset is MMS capable is not specified by the standards. A database is usually maintained by the operator, and in it each mobile phone number is marked as being associated with a legacy handset or not. It can be a bit hit and miss since customers can change their handset at will and this database is not usually updated dynamically. E-mail and web-based gateways to the MMS (and SMS) system are common. On the reception side, the content servers can typically receive service requests both from WAP and normal HTTP browsers, so delivery via the web is simple. For sending from external sources to handsets, most carriers allow MIME encoded message to be sent to the receiver's phone number with a special domain.

Challenges

There are some interesting challenges with MMS that do not exist with SMS:

- Content adaptation: Multimedia content created by one brand of MMS phone may not be entirely compatible with the capabilities of the recipient's MMS phone. In the MMS architecture, the recipient MMSC is responsible for providing for *content adaptation* (e.g., image resizing, audio codec transcoding, etc.), if this feature is enabled by the mobile network operator. When content adaptation is supported by a network operator, its MMS subscribers enjoy compatibility with a larger network of MMS users than would otherwise be available.

- Distribution lists: Current MMS specifications do not include distribution lists nor methods by which large numbers of recipients can be conveniently addressed, particularly by content providers, called *Value-added service providers* (VASPs) in 3GPP. Since most SMSC vendors have adopted FTP as an ad-hoc method by which large distribution lists are transferred to the SMSC prior to being used in a bulk-messaging SMS submission, it is expected that MMSC vendors will also adopt FTP.
- Bulk messaging: The flow of *peer-to-peer* MMS messaging involves several over-the-air transactions that become inefficient when MMS is used to send messages to large numbers of subscribers, as is typically the case for VASPs. For example, when one MMS message is submitted to a very large number of recipients, it is possible to receive a *delivery report* and *read-reply report* for each and every recipient. Future MMS specification work is likely to optimize and reduce the transactional overhead for the bulk-messaging case.

- Handset Configuration: Unlike SMS, MMS requires a number of handset parameters to be set. Poor handset configuration is often blamed as the first point of failure for many users. Service settings are sometimes preconfigured on the handset, but mobile operators are now looking at new device management technologies as a means of delivering the necessary settings for data services (MMS, WAP, etc.) via over-the-air programming (OTA).
- WAP Push: Few mobile network operators offer direct connectivity to their MMSCs for content providers. This has resulted in many content providers using WAP push as the only method available to deliver 'rich content' to mobile handsets. WAP push enables 'rich content' to be delivered to a handset by specifying the URL (via binary SMS) of a pre-compiled MMS, hosted on a content provider's web server. A consequence is that the receiver who pays WAP per kb or minute (as opposed to a flat monthly fee) pays for receiving the MMS, as opposed to only paying for sending one, and also paying a different rate.

Although the standard does not specify a maximum size for a message, 300 kB is the current recommended size used by networks due to some limitations on the WAP gateway side.

Interfaces

- MM1: the 3GPP interface between MMS User Agent and MMS Center
- MM2: the 3GPP interface between MMS Relay and MMS Server
- MM3: the 3GPP interface between MMS Center and external servers
- MM4: the 3GPP interface between MMS Centers
- MM5: the 3GPP interface between MMS Center and HLR
- MM6: the 3GPP interface between MMS Center and user databases
- MM7: the 3GPP interface between MMS VAS applications and MMS Center
- MM8: the 3GPP interface between MMS Center and the billing systems
- MM9: the 3GPP interface between MMS Center and an online charging system
- MM10: the 3GPP interface between MMS Center and a message service control function
- MM11: the 3GPP interface between MMS Center and an external transcoder

MULTIMEDIA WEB ONTOLOGY LANGUAGE

Multimedia Web Ontology Language (MOWL) has been designed to facilitate semantic interactions with multimedia contents. It supports perceptual modeling of concepts using expected media properties. While the reasoning in traditional ontology languages, e.g. Web Ontology Language (OWL), is based on Description Logics, MOWL supports a probabilistic reasoning framework based on Bayesian Network.

History

W3C forum has undertaken the initiative of standardizing the ontology representation for web-based applications. The Web Ontology Language (OWL), standardized in 2004 after maturing through XML(S), RDF(S) and DAML+OIL is a result of that effort. Ontology in OWL (and some of its predecessor languages) has been successfully used in establishing semantics of text in specific application contexts. The concepts and properties in these traditional ontology languages are expressed as text, making an ontology readily usable for semantic analysis of textual documents. Semantic processing of media data calls for perceptual modeling of domain concepts with their media properties. Such modeling was first proposed in the Ph.D. Thesis by Hiranmay Ghosh (Electrical Engineering Department, IIT Delhi, 2002) in the form of *Knowledge Description Language* (KDL). With the standardization of OWL by W3C, KDL was merged with OWL to form Multimedia Web Ontology Language (MOWL). Several students have taken the work further to implement research prototypes of retrieval systems and ontology learning.

Key Features

Syntactically, MOWL is an extension of OWL. These extensions enable

- Definition of media properties following MPEG-7 media description model.
- Probabilistic association of media properties with the domain concepts.
- Formal semantics to the media properties to enable reasoning.
- Formal semantics for spatio-temporal relations across media objects and events.

MOWL is accompanied with reasoning tools that support

- Construction of model of observation for a concept in multimedia documents with expected media properties.
- Probabilistic (Bayesian) reasoning for concept recognition with the model of observation.

INTERNET TELEVISION

Internet television (otherwise known as Internet TV, or Online TV, and not to be confused with Web television or IPTV), is a television service distributed via the Internet. Some Internet television is known as Catch up TV. It has become very popular with services such as BBC iPlayer, 4oD, ITV Player (also STV Player and UTV Player) and Demand Five in the United Kingdom, Hulu and Revision3 in the United States, Nederland 24 in the Netherlands, ABC iView and Australia

Live TV in Australia, and SeeSaw, RTÉ Player in the Republic Of Ireland, Tivibu in Turkey ; see List of Internet television providers.

Concept

Internet television allows the users to choose the programme or the TV show they want to watch from an archive of programmes or from a channel directory. The two forms of viewing Internet television are streaming the content directly to a media player or simply downloading the programme to the user's computer. With the "TV on Demand" market growing, these on-demand websites or applications are a must have for major television broadcasters. For example the BBC iPlayer brings in users which stream more than one million videos per week, with one of the BBC's headline shows "The Apprentice" taking over 3 - 5% of the UK's internet traffic due to people watching the first episode on the BBC iPlayer. Every night the use of on-demand TV peaks at around 10 pm, Most providers of the service provide several different formats and quality controls so that the service can be viewed on many different devices.

Some services now offer a HD service along side their SD, streaming is the same but offers the quality of HD to the device being used, as long as it is using a HD screen. During Peak times the BBC iPlayer transmits 12 GB (gigabytes) of information per second. Over the course of a month the iPlayer sends 7 PB (petabytes) of information. Before 2006, most Catch-up services used peer-to-peer (P2P) networking, in which users downloaded an application and data would be shared between the users rather than the service provider giving the now more commonly used streaming method. Now most service providers have moved away from the P2P systems and are now using the streaming media. This is good for the service provider as in the old P2P system the distribution costs were high and the servers normally couldn't handle the large amount of downloading and data transfer.

Market Competitors

Many providers of internet television services exist including conventional television stations that have taken advantage of the internet as way to continue showing programmes after they have been broadcast often advertised as "On-demand" and "Catch-up" services. Today, almost every major broadcaster around the world is operating an internet television platform. Examples include the BBC, which introduced the BBC iPlayer on 25 June 2008 as an extension to its "RadioPlayer" and already existing streamed video clip content, and Channel 4 that launched 4oD ("4 on Demand") in November 2006 allowing users to watch recently shown content. Most internet television services allow users to view content free of charge however some content is charged for. Other internet television providers include Australia Live TV SeeSaw, ITV player, TVCatchup, Demand Five, Eurosport player and Sky Player.

Access/Usability

The ability to access internet television is heavily dependent on internet streaming speeds. This limits adoption in many countries, as broadband penetration is limited; in the European Union only 25% of consumers had access to Broadband internet in 2010. Using an Internet Service Provider, something which is common in many homes in the developed world, the user simply enters their chosen website address. For example, bbc.co.uk/iplayer or http://video.pbs.org . If the user has no select preference of streaming service, the name of a chosen television programme can be inputted into a search engine followed by a phrase such as "online streaming" or "watch on the net". Accessing television on the internet has never been so simple, due to this usability of streaming services has had to be improved to maintain the simplicity of the process. Upon selection of a programme and website, the user may have to wait a few seconds or minutes to allow their desired programme to stream. A process called buffering allows the programme to run in one smooth showing as opposed to stopping and starting to allow the programme to stream.

Control

Controlling content on the Internet presents a challenge for most providers; to try to ensure that a user is allowed to view content such as programmes with age certificates, providers use methods such as parental controls that allows restrictions to be placed upon the use and access of certificated material. The BBC iPlayer makes use of a parental control system giving parents the option to "lock" content, meaning that a password would have to be used to access it. Flagging systems can be used to warn a user that content may be certified or that it may be post watershed for a programme. Honor systems are also used where users are asked for their dates of birth or age to verify if they are able to view certain content.

Archives

An archive is a collection of information and media much like a library or interactive storage facility. It is a necessity for an on-demand media service to maintain archives so that users can watch programmes that have already been aired on standard broadcast television. However, these archives can vary from a few weeks to months to years, depending on the curator and what programme it is. In contrast 4oD channel 4's on-demand service offers many of its much older programmes as well that were originally aired years ago. An example of this is the comedy *The IT Crowd* where users can view the full series on the internet player. The same is true for other hit channel 4 comedies such as *The Inbetweeners* and *Black Books*. Having an extensive archive however can bring problems along with benefits. Large archives are expensive to maintain, server farms and mass storage is needed along with ample bandwidth to transmit it all. Vast archives can be hard to catalogue and sort so that it is accessible to users. The benefits in most cases outweigh these problems. This is because large archives bring in far more users who in turn watch more media, leading to a wider audience base and more advertising revenue. Large archives will also mean the user will spend more time on that website rather than a competitors, leading to starvation of demand for the competitors.

Broadcasting Rights

Broadcasting rights change from country to country and even within provinces of countries. These rights govern the distribution of copyrighted content and media and allow the sole distribution of that content at any one time. An example of programmes only being aired in certain countries is BBC iPlayer. Users can only stream content from the BBC iPlayer from Britain because the BBC only allows free use of their product for users within the United Kingdom because those users pay a TV license to fund part of the BBC. Broadcasting rights can also be restricted to allowing a broadcaster rights to distribute that content for a limited time. Channel 4's online service 4oD can only stream shows created in the US by companies such as "HBO" for 30 days after they are aired on one of the Channel 4 group channels. This is to boost DVD sales for the companies who produce that media. Some companies pay very large amounts for broadcasting rights with sports and US sitcoms usually fetching the highest price from UK based broadcasters.

Profits and Costs

With the exception of Internet connectivity costs many online television channels or sites are free. These sites maintain this free TV policy through the use of video advertising, short commercials and banner adverts may show up before a video is played. An example of this is on the abc.com catch up website; in place of the advert breaks on normal television a short 30 second advert is played. This short advertising time means that the user does not get fed up and money can be made off of advertising, to allow web designers to offer quality content which would otherwise cost. This is how online TV makes a profit.

Technologies used for Internet Television

The Hybrid Broadcast Broadband TV (HbbTV) consortium of industry companies (such as SES Astra, Humax, Philips, and ANT Software) is currently promoting and establishing an open European standard (called HbbTV) for hybrid set-top boxes for the reception of broadcast and broadband digital TV and multimedia applications with a single user interface. Current providers of internet television use various technologies to provide a service such as peer-to-peer (P2P) technologies, VoD systems, and live streaming. BBC iPlayer makes use of Adobe Flash Player to provide streaming video clips and other software provided by Adobe for its download service. NBC, Bloomberg Television, and Showtime use live streaming services from BitGravity to stream live television to paid subscribers using a standard http protocol. DRM (digital rights management) software is also incorporated into many internet television services Sky Player has software that is provided by Microsoft to prevent content being copied. Internet television is also cross platform, the Sky Player service has been expanded to the Xbox 360 on October 27 and to Windows Media Center and then to Windows 7 PCs on November 19.

BBC iPlayer is also available through Virgin Media's on-demand service and other platforms such as FetchTV and games consoles including the Xbox 360, Wii and the PlayStation 3. Other platforms that internet television is available on include mobile platforms such as the iPhone and iPod Touch, Nokia N96, Sony Ericsson C905 and many other mobile devices.

Website vs. Applications

The main problem with on-demand video services that are applications on desktop computers is getting users to download them and register. It is far easier for a user to simply log onto a webpage without registering than to have to spend time registering and downloading often large programmes.

However applications are more powerful in that they can manage the downloading of content far better and these programmes can usually be watched offline for 30 days after downloading.

Stream Quality

Stream quality refers to the quality of the image and audio transferred from the servers of the distributor to the home screen on a user. Higher quality video such as video in high definition (720p+) requires higher bandwidth and faster connection speeds. The general accepted kbps download rate needed to stream high definition video that has been encoded with H.264 is 3500, where as standard definition TV can range from 500 to 1500 kbps depending on the resolution on screen. In the UK, the BBC iPlayer deals with the largest amount of traffic yet it offers HD content along with SD content. As more people get internet connections which can deal with streaming HD video over the internet the BBC iPlayer has tried to keep up with demand and pace. However, as streaming HD video takes around 1.5gb of data per hour of video it took a lot of investment by the BBC to implement this on such a large scale. For users which do not have the bandwidth to stream HD video or even high SD video which requires 1500kbps, the BBC iPlayer offers lower bitrate streams which in turn leads to lower video quality. This makes use of an adaptive bitrate stream so that if the users bandwidth suddenly drops, iPlayer will lower its streaming rate to compensate for this. This diagnostic tool offered on the BBC iPlayer site measures a user's streaming capabilities and bandwidth for free. Although competitors in the UK such as 4oD, ITV Player and Demand Five have not yet offered

HD streaming, the technology to support it is fairly new and widespread HD streaming is not an impossibility. The availability of Channel 4 and Five programmes on YouTube is predicted to prove incredibly popular as series such as *Skins*, *Green Wing*, *The X Factor* and others become available in a simple, straightforward format on a website which already attracts millions of people every day.

MULTIMEDIA SEARCH

Multimedia search enables information search using search queries in multiple data types including text and other multimedia formats. Multimedia search can be implemented through multimodal search interfaces, i.e., interfaces that allow to submit search queries not only as textual requests, but also through other media. This is often referred to as Search by example, because the typical interaction consists in submitting a piece of information (e.g., a video, an image, or a piece of audio) at the purpose of finding similar multimedia items.

Multimedia search includes:

- Audio search
- Visual search

Audio search is a search where you can search for songs and music. Visual search is a search where you can search for videos and pictures.

The most major search engines give the user the opportunity to search for audio, video and pictures. However, most of them only allow to search with a mono-modal interaction, based on textual keyword search.

IP MULTIMEDIA SUBSYSTEM

The IP Multimedia Subsystem (IMS) is an architectural framework for delivering Internet Protocol (IP) multimedia services. It was originally designed by the wireless standards body 3rd Generation Partnership Project (3GPP), as a part of the vision for evolving mobile networks beyond GSM. Its original formulation (3GPP R5) represented an approach to delivering "Internet services" over GPRS. This vision was later updated by 3GPP, 3GPP2 and TISPAN by requiring support of networks other than GPRS, such as Wireless LAN, CDMA2000 and fixed line. To ease the integration with the Internet, IMS uses IETF protocols wherever possible, e.g. Session Initiation Protocol (SIP). According to the 3GPP, IMS is not intended to standardize applications but rather to aid the access of multimedia and voice applications from wireless and wireline terminals, i.e. create a form of fixed-mobile convergence (FMC). This is done by having a horizontal control layer that isolates the access network from the service layer. From a logical architecture perspective, services need not have their own control functions, as the control layer is a common horizontal layer. However in implementation this does not necessarily map into greater reduced cost and complexity. Alternative and overlapping technologies for access and provisioning of services across wired and wireless networks include combinations of Generic Access Network, soft switches and “naked” SIP. It is easier to sell services than to sell the virtues of “integrated services”, but additionally the task to sell an IMS based on a single service is also difficult as there are often (cheaper) alternatives to creating and deploying that particular service. Since it is becoming increasingly easier to access content and contacts using mechanisms outside the control of traditional wireless/fixed operators, the interest of IMS is being challenged.

History

- IMS was originally defined by an industry forum called 3G.IP, formed in 1999. 3G.IP developed the initial IMS architecture, which was brought to the 3rd Generation Partnership Project (3GPP), as part of their standardization work for 3G mobile phone systems in UMTS networks. It first appeared in Release 5 (evolution from 2G to 3G networks), when SIP-based multimedia was added. Support for the older GSM and GPRS networks was also provided.
- 3GPP2 (a different organization from 3GPP) based their CDMA2000 Multimedia Domain (MMD) on 3GPP IMS, adding support for CDMA2000.
- 3GPP release 6 added interworking with WLAN, Inter-operability between IMS using different IP-connectivity networks, routing group identities, multiple

registration and forking, presence, speech recognition and speech-enabled services (Push to talk).

- 3GPP release 7 added support for fixed networks, by working together with TISPAN release R1.1, the function of AGCF (Access Gateway control function) and PES (PSTN Emulation Service) are introduced to the wire-line network for the sake of inheritance of services which can be provided in PSTN network. Also added voice call continuity between circuit switching and packet switching domain (VCC), fixed broadband connection to the IMS, interworking with non-IMS networks, Policy and Charging Control (PCC), emergency sessions.
- 3GPP release 8 added support for Long Term Evolution (LTE), System Architecture Evolution (SAE), Multimedia Session Continuity, Enhanced emergency sessions and IMS centralized services.

The IP Multimedia Core Network Subsystem is a collection of different functions, linked by standardized interfaces, which grouped form one IMS administrative network. A function is not a node (hardware box): an implementer is free to combine 2 functions in 1 node, or to split a single function into 2 or more nodes.

Each node can also be present multiple times in a single network, for dimensioning, load balancing or organizational issues.

Access Network

The user can connect to an IMS network in various ways, most of which use the standard Internet Protocol (IP). IMS terminals (such as mobile phones, personal digital assistants (PDAs) and computers) can register directly on an IMS network, even when they are roaming in another network or country (the visited network). The only requirement is that they can use IP and run Session Initiation Protocol (SIP) user agents.

Fixed access (e.g., Digital Subscriber Line (DSL), cable modems, Ethernet), mobile access (e.g. W-CDMA, CDMA2000, GSM, GPRS) and wireless access (e.g. WLAN, WiMAX) are all supported. Other phone systems like plain old telephone service (POTS—the old analogue telephones), H.323 and non IMS-compatible VoIP systems, are supported through gateways.

Core Network

Home subscriber server: The *Home Subscriber Server* (HSS), or *User Profile Server Function* (UPSF), is a master user database that supports the IMS network entities that actually handle calls. It contains the subscription-related information (subscriber profiles), performs authentication and authorization of the user, and can provide information about the subscriber's location and IP information. It is similar to the GSM Home Location Register (HLR) and Authentication Centre (AuC).

A *Subscriber Location Function* (SLF) is needed to map user addresses when multiple HSSs are used.

User identities: Various identities may be associated with IMS: IP Multimedia Private Identity (IMPI), IP Multimedia Public Identity (IMPU), Globally Routable User Agent URI (GRUU), Wildcarded Public User Identity. Both IMPI and IMPU are not phone numbers or other series of digits, but Uniform Resource Identifier (URIs), that can be digits.

IP Multimedia Private Identity

The *IP Multimedia Private Identity* (IMPI) is a unique permanently allocated global identity assigned by the home network operator, and is used, for example, for Registration, Authorization, Administration, and Accounting purposes. Every IMS user shall have one or more IMPI.

IP Multimedia Public Identity

The *IP Multimedia Public Identity* (IMPU) is used by any user for requesting communications to other users (e.g. this might be included on a business card). There can be multiple IMPU per IMPI. The IMPU can also be shared with another phone, so that both can be reached with the same identity (for example, a single phone-number for an entire family).

Globally Routable User Agent URI

Globally Routable User Agent URI (GRUU) is an identity that identifies a unique combination of IMPU and UE instance. There are two types of GRUU: Public-GRUU (P-GRUU) and Temporary GRUU (T-GRUU).

- P-GRUU reveal the IMPU and are very long lived.
- T-GRUU do not reveal the IMPU and are valid until the contact is explicitly de-registered or the current registration expires

Wild Carded Public User Identity

A *wildcarded Public User Identity* expresses a set of IMPU grouped together. The HSS subscriber database contains, the IMPU, IMPI, IMSI, and MSISDN, subscriber service profiles, service triggers and other information.

Call/Session Control

Several roles of Session Initiation Protocol (SIP) servers or proxies, collectively called Call Session Control Function (CSCF), are used to process SIP signalling packets in the IMS.

- A *Proxy-CSCF* (P-CSCF) is a SIP proxy that is the first point of contact for the IMS terminal. It can be located either in the visited network (in full IMS networks) or in the home network (when the visited network is not IMS compliant yet). Some networks may use a Session Border Controller for this function. The P-CSCF is at its core a specialized SBC for the User–network interface which not only protects the network, but also the IMS terminal. The use of additional SBC between the IMS terminal and the P-CSCF as such pointless and also not feasible due to the signaling being encrypted on this leg. The terminal discovers its P-CSCF with either DHCP, or it may be configured (e.g. during initial provisioning or via a 3GPP IMS Management Object (MO)) or in the ISIM or assigned in the PDP Context (in General Packet Radio Service (GPRS)).

- o it is assigned to an IMS terminal before registration, and does not change for the duration of the registration
- o it sits on the path of all signalling, and can inspect every signal; the IMS terminal must ignore any other unencrypted signalling
- o it provides subscriber authentication and may establish an IPsec or TLS security association with the IMS terminal. This prevents spoofing attacks and replay attacks and protects the privacy of the subscriber.
- o it inspects the signaling and ensures that the IMS terminals do not misbehave (e.g. change normal signaling routes, do not obey home network's routing policy)
- o it can also compress and decompress SIP messages using SigComp, which reduces the round-trip over slow radio links
- o it may include a Policy Decision Function (PDF), which authorizes media plane resources e.g. quality of service (QoS) over the media plane. It is

used for policy control, bandwidth management, etc. The PDF can also be a separate function.

- o it also generates charging records

- A *Serving-CSCF* (S-CSCF) is the central node of the signalling plane. It is a SIP server, but performs session control too. It is always located in the home network. It uses Diameter Cx and Dx interfaces to the HSS to download user profiles and upload user-to-S-CSCF associations (the user profile is only cached locally for processing reasons only and is not changed). All necessary subscriber profile information is loaded from the HSS.
  - o it handles SIP registrations, which allows it to bind the user location (e.g. the IP address of the terminal) and the SIP address
  - o it sits on the path of all signaling messages of the locally registered users, and can inspect every message
  - o it decides to which application server(s) the SIP message will be forwarded, in order to provide their services
  - o it provides routing services, typically using Electronic Numbering (ENUM) lookups
  - o it enforces the policy of the network operator
  - o there can be multiple S-CSCFs in the network for load distribution and high availability reasons. It's the HSS that assigns the S-CSCF to a user, when it's queried by the I-CSCF. There are multiple options for this purpose, including a

mandatory/optional capabilities to be matched between subscribers and S-CSCFs.

- An *Interrogating-CSCF* (I-CSCF) is another SIP function located at the edge of an administrative domain. Its IP address is published in the Domain Name System (DNS) of the domain (using NAPTR and SRV type of DNS records), so that remote servers can find it, and use it as a forwarding point (e.g. registering) for SIP packets to this domain.
  - o it queries the HSS to retrieve the address of the S-CSCF and assign it to a user performing SIP registration
  - o it also forwards SIP request or response to the S-CSCF
  - o Up to Release 6 it can also be used to hide the internal network from the outside world (encrypting parts of the SIP message), in which case it's called a *Topology Hiding Inter-network Gateway* (THIG). From Release 7 onwards this "entry point" function is removed from the I-CSCF and is now part of the *Interconnection Border Control Function* (IBCF). The IBCF is used as gateway to external networks, and provides NAT and Firewall functions (pinholing). The IBCF is practically a Session Border Controller specialized for the NNI.

Application Servers

SIP Application servers (AS) host and execute services, and interface with the S-CSCF using Session Initiation Protocol

(SIP). An example of an application server that is being developed in 3GPP is the Voice call continuity Function (VCC Server). Depending on the actual service, the AS can operate in SIP proxy mode, SIP UA (user agent) mode or SIP B2BUA mode. An AS can be located in the home network or in an external third-party network. If located in the home network, it can query the HSS with the Diameter Sh or Si interfaces (for a SIP-AS).

- SIP AS: Host and execute IMS specific services
- *IP Multimedia Service Switching Function* (IM-SSF): Interfaces SIP to CAP to communicate with CAMEL Application Servers
- OSA Service Capability Server (OSA SCS) : Interfaces SIP to the OSA framework

Functional model: The AS-ILCM and AS-OLCM store transaction state, and may optionally store session state depending on the specific service being executed. The AS-ILCM interfaces to the S-CSCF (ILCM) for an incoming leg and the AS-OLCM interfaces to the S-CSCF (OLCM) for an outgoing leg. Application Logic provides the service(s) and interacts between the AS-ILCM and AS-OLCM.

Public Service Identity: Public Service Identities (PSI) are identities that identify services, which are hosted by Application Servers. As user identities, PSI shall take the form of either a SIP or Tel URI. PSIs are stored in the HSS either as a distinct PSI or as a wildcarded PSI:

- a distinct PSI contains the PSI that is used in routing
- a wildcarded PSI represents a collection of PSIs.

Media Servers

The *Media Resource Function* (MRF) provides media related functions such as media manipulation (e.g. voice stream mixing) and playing of tones and announcements.

Each MRF is further divided into a *Media Resource Function Controller* (MRFC) and a *Media Resource Function Processor* (MRFP).

- The MRFC is a signalling plane node that interpret information coming from an AS and S-CSCF to control the MRFP
- The MRFP is a media plane node used to mix, source or process media streams. It can also manage access right to shared resources.

The *Media Resource Broker* (MRB) is a functional entity that is responsible for both collection of appropriate published MRF information and supplying of appropriate MRF information to consuming entities such as the AS. MRB can be used in two modes:

- Query mode: AS queries the MRB for media and sets up the call using the response of MRB
- In-Line Mode: AS sends a SIP INVITE to the MRB. The MRB sets up the call

Breakout Gateway

A *Breakout Gateway Control Function* (BGCF) is a SIP proxy which processes requests for routing from an S-CSCF when the S-CSCF has determined that the session cannot be routed using DNS or ENUM/DNS. It includes routing functionality based on telephone numbers.

PSTN Gateways

A PSTN/CS gateway interfaces with PSTN circuit switched (CS) networks. For signalling, CS networks use ISDN User Part (ISUP) (or BICC) over Message Transfer Part (MTP), while IMS uses Session Initiation Protocol (SIP) over IP. For media, CS networks use Pulse-code modulation (PCM), while IMS uses Real-time Transport Protocol (RTP).

- A Signalling Gateway (SGW) interfaces with the signalling plane of the CS. It transforms lower layer protocols as Stream Control Transmission Protocol (SCTP, an Internet Protocol (IP) protocol) into Message Transfer Part (MTP, an Signalling System 7 (SS7) protocol), to pass ISDN User Part (ISUP) from the MGCF to the CS network.
- A *Media Gateway Controller Function* (MGCF) is SIP endpoint that does call control protocol conversion between SIP and ISUP/BICC and interfaces with the SGW over SCTP. It also controls the resources in a *Media Gateway* (MGW) across an H.248 interface.

- A *Media Gateway* (MGW) interfaces with the media plane of the CS network, by converting between RTP and PCM. It can also transcode when the codecs don't match (e.g. IMS might use AMR, PSTN might use G.711).

Media Resources

Media Resources are those components that operate on the media plane and are under the control of IMS Core functions. Specifically, *Media Server* (MS) and *Media gateway* (MGW).

NGN Interconnection

There are two types of Next Generation Networking Interconnection:

- *Service oriented Interconnection* (SoIx): The physical and logical linking of NGN domains that allows carriers and service providers to offer services over NGN (i.e. IMS and PES) platforms with control, signalling (i.e. session based), which provides defined levels of interoperability. For instance, this is the case of "carrier grade" voice and/or multimedia services over IP interconnection. "Defined levels of interoperability" are dependent upon the service or the QoS or the Security, etc.
- *Connectivity oriented Interconnection* (CoIx): The physical and logical linking of carriers and service

providers based on simple IP connectivity irrespective of the levels of interoperability. For example, an IP interconnection of this type is not aware of the specific end to end service and, as a consequence, service specific network performance, QoS and security requirements are not necessarily assured. This definition does not exclude that some services may provide a defined level of interoperability. However only SoIx fully satisfies NGN interoperability requirements.

An NGN interconnection mode can be direct or indirect. Direct interconnection refers to the interconnection between two network domains without any intermediate network domain.

Indirect interconnection at one layer refers to the interconnection between two network domains with one or more intermediate network domain(s) acting as transit networks. The intermediate network domain(s) provide(s) transit functionality to the two other network domains. Different interconnection modes may be used for carrying service layer signalling and media traffic.

Charging

Offline charging is applied to users who pay for their services periodically (e.g., at the end of the month). Online charging, also known as credit-based charging, is used for prepaid services, or real-time credit control of postpaid services. Both may be applied to the same session.

Charging function addresses are addresses distributed to each IMS entities and provide a common location for each entity to send charging information. *Charging Data Function* (CDF) addresses are used for offline billing and *Online Charging Function* (OCF) for online billing.

- Offline Charging : All the SIP network entities (P-CSCF, I-CSCF, S-CSCF, BGCF, MRFC, MGCF, AS) involved in the session use the Diameter Rf interface to send accounting information to a CDF located in the same domain. The CCF will collect all this information, and build a *Call Detail Record* (CDR), which is sent to the billing system (BS) of the domain.

IMS Charging Identifier (ICID) as a unique identifier generated by the first IMS entity involved in a SIP transaction and used for the correlation with CDRs. *Inter Operator Identifier* (IOI) is a globally unique identifier shared between sending and receiving networks. Each domain has its own charging network. Billing systems in different domains will also exchange information, so that roaming charges can be applied.

- Online charging : The S-CSCF talks to a *Session Charging Function* (SCF) which looks like a regular SIP application server. The SCF can signal the S-CSCF to terminate the session when the user runs out of credits during a session. The AS and MRFC use the Diameter Ro interface towards an OCF.

- o When *Immediate Event Charging* (IEC) is used, a number of credit units is immediately deducted from the user's account by the ECF and the MRFC or AS is then authorized to provide the service. The service is not authorized when not enough credit units are available.
- o When *Event Charging with Unit Reservation* (ECUR) is used, the ECF first reserves a number of credit units in the user's account and then authorizes the MRFC or the AS. After the service is over, the number of spent credit units is reported and deducted from the account; the reserved credit units are then cleared.

Session Handling

One of the most important features of IMS, that of allowing for a SIP application to be dynamically and differentially (based on the user's profile) triggered, is implemented as a filter-and-redirect signalling mechanism in the S-CSCF. The S-CSCF might apply filter criteria to determine the need to forward SIP requests to AS. It is important to note that services for the originating party will be applied in the originating network, while the services for the terminating party will be applied in the terminating network, all in the respective S-CSCFs.

Initial Filter Criteria

Initial Filter Criteria (IFC) are filter criteria that are stored in the HSS as part of the IMS Subscription Profile and are downloaded to the S-CSCF upon user registration (for registered users) or on processing demand (for services, acting as unregistered users). They represent a provisioned subscription of a user to an application. iFC are valid throughout the registration lifetime or until the User Profile is changed. The term Shared iFC denotes an iFC which, due to its common use for a large number of subscribers, is only referenced in the Subscription Profile and provisioned on a different path between the HSS and the S-CSCF.

The iFC may be composed of:

- An Application Server URI where the request is to be forwarded in case of a match.
- A Trigger Point in the form of a logical condition which is verified against initial dialog creating SIP requests or stand-alone SIP requests.

Security Aspects of Early IMS and Non-3GPP Systems

It is envisaged that security defined in TS 33.203 may not be available for a while especially because of the lack of USIM/ ISIM interfaces and prevalence of devices that support IPv4. For this situation, to provide some protection against the most significant threats, 3GPP defines some security mechanisms, which are informally known as “early IMS security,” in TR33.978. This mechanism relies on the authentication performed during the network attachment procedures, which binds between the user’s profile and its IP address. This mechanism is also weak because the signaling is not protected on the User–network interface. CableLabs in PacketCable#PacketCable_2.0, which adopted also the IMS architecture but has no USIM/ISIM capabilities in their terminals, published deltas to the 3GPP specifications where the Digest-MD5 is a valid authentication option. Later on, also TISPAN did a similar effort given their Fixed Networks scopes, yet the procedures are different. To compensate for the lack of IPsec capabilities, TLS has been added as an option for securing the Gm interface. Later 3GPP Releases have included the Digest-MD5 method, towards a Common-IMS platform, yet in its own and again different approach.

Although all 3 variants of Digest-MD5 authentication have the same functionality and are the same from the IMS terminal's perspective, the implementations on the Cx interface between the S-CSCF and the HSS are different.

MOTION GRAPHIC DESIGN

Motion Design is a subset of graphic design in that it uses graphic design principles in a film or video context (or other temporally evolving visual medium) through the use of animation or filmic techniques. Examples include the typography and graphics you see as the titles for a film, or opening sequences for television or the spinning, web-based animations, three-dimensional logo for a television channel. About 12 minutes in every hour of broadcast television is the work of the motion graphics designer, yet it is known as the invisible art, as many viewers are unaware of this component of programming. Although this art form has been around for decades, it has taken quantum leaps forward in recent years, in terms of technical sophistication. If you watch much TV or see many films, you will have noticed that the graphics, the typography, and the visual effects within this medium have become much more elaborate and sophisticated.

Technology

The elevation of this art form is largely due to technology improvements. Computer programmes for the film and video industry have become vastly more powerful and more available. Probably the leading programme used by motion graphic designers is Adobe After Effects, which allows them to create and modify graphics over time. Adobe After Effects is sometimes referred to as "Photoshop for film." A relatively recent product in the market is Apple Inc. Motion, now a part of Final Cut Studio. Adobe Flash is widely used to create motion design for the web. A typical motion designer is a person trained in traditional graphic design who has learned to integrate the elements of time, sound and space into his/her existing skill-set of design knowledge. Motion designers can also come from filmmaking or animation backgrounds.

Notable Motion Designers

- Saul Bass
- Maurice Binder
- Pablo Ferro

Motion Design & Digital Compositing Software Packages

Since motion design is created using images and video sequences, a great complementary tool is a 3d software package. Cinema 4D is widely used for its intuitive interface, layered export to Adobe After Effects, and the additional MoGraph module, but there are also several others. Such packages can generate images or video sequences with an alpha channel, which stores all the transparency information.

Motion Design applications include:

- Adobe After Effects
- Jahshaka
- Autodesk Combustion
- Apple Motion/Shake
- Max/MSP
- Apple Quartz Composer
- Various VJ Programmes
- Smith Micro Software Anime Studio
- Adobe Flash

3D Programmes used in motion graphics include:

- Maxon Cinema 4D
- Softimage XSI
- Autodesk 3d studio max
- Autodesk Maya
- NewTek Lightwave
- e-on Vue Infinite

- The Blender Foundation Blender software
- EI Technology Group Electric Image Animation System

Motion graphics plugins include:

- Magic bullet

MULTIMEDIA ARTISTS

Multimedia artists are contemporary artists who use a wide range of media to communicate their art. Such media range from installation art, to rooms containing found objects or other material, to kinetic sculpture, to sound and visual effects. It is important to distinguish between multimedia art and mixed media artworks. Within the visual arts, *mixed media* tends to refer to work that combines various traditionally distinct visual art media - such as certain works of Frank Stella or Jane Frank which merge painting and sculpture, for example. A work on canvas that combines oil paint, newspaper collage, chalk, glass, and ink, for example, could be called a "mixed media" work - but not a work of "multimedia art." *Multimedia* art implies a broader scope than *mixed media*, as in creations combining visual art media with elements usually considered the proper domain of (for example) literature, drama, dance, filmmaking, or music. Multimedia artwork also frequently engages senses other than sight, such as hearing, touch, or smell. A multimedia artwork can also move, occupy time, or develop over a span of time, instead of remaining static (as does a traditional painting or sculpture). Another frequent trait of multimedia artworks is the use of advanced technological means, such as electronic or computer-generated sound, video, animation, and interactivity. Certain traditional genres such as opera and film are inherently multidisciplinary or even "multimedia" in a very loose sense, since they involve drama, literature, visual art, music, dance, and costumes. Indeed, a union of the arts was exactly what Richard Wagner imagined in his ideal of the "Gesamtkunstwerk" or a "synthesis of the arts" (literally: "complete artwork"). Nevertheless, in contemporary terms, opera or even movies would not properly be considered "multimedia art." A work of multimedia art is usually on a smaller scale than an opera or a movie, much less tradition-bound, and typically created entirely by a single person (rather than the collaborative effort of opera or moviemaking). A multimedia work also usually does not require performers. If human performers are used, they are usually ordinary, untrained people, doing nothing requiring any advanced or traditional training, as opposed to trained singers or actors. Multimedia artwork is often presented in a curated museum or gallery setting, in which the piece is understood to be an extended form of visual art. The creator of a multimedia work of art is typically someone with a formal background in visual art.

MULTIMEDIA LEARNING

Multimedia learning is the common name used to describe the cognitive theory of multimedia learning. This theory encompasses several principles of learning with multimedia.

The Modality Principle

When information is in fact better remembered when accompanied by a visual image. Baddeley and Hitch proposed a theory of working memory in 1974 which has two largely independent subcomponents that tend to work in parallel - one visual and one verbal/acoustic. This allows us to simultaneously process information coming from our eyes and ears. Thus a learner is not necessarily overwhelmed or overloaded by multimodal instruction, and it can in fact be beneficial. The finding that items presented both visually and verbally are better remembered gave rise to dual-coding theory, first proposed by Paivio and later applied to multimedia by Richard Mayer and his associates. Mayer has shown learners are better able to transfer their learning given multimodal instruction. Mayer explains the modality effect from an information processing/cognitive load perspective. In a series of studies Mayer and his colleagues tested Paivio's dual-coding theory, with multimedia. They repeatedly found that students learning given multimedia with animation and narration consistently did better on transfer questions than those who learn from animation and text-based materials. That is, they were significantly better when it came to applying what they had learned after receiving multimedia rather than mono-media (visual only) instruction. These results were then later confirmed by other groups of researchers. Initially the instructional content of these multimedia learning studies was limited to logical scientific processes that centered on cause-and-effect systems like automobile braking systems, how a bicycle pump works, or cloud formation. But eventually it was found that the modality effect could be extended to other domains, which were not necessarily cause-and-effect based systems. Information then can and should be encoded as both as visually and auditory (narration). If verbal information is encoded auditorily it reduces the cognitive load of the learner and they are better able to handle that incoming information. Mayer has since called this the "Modality effect," or the Modality Principle. This was one of the many principles of his "Cognitive Theory of Multimedia Learning".

The Redundancy Principle

According to this principle: "Students learn better from animation and narration than from animation, narration, and on-screen text." Thus it's better to eliminate redundant material. This is because learners do not learn as well when they both hear and see the same verbal message during a presentation. This is a special case of the split attention effect of Sweller and Chandler.

Other Principles

- Spatial Contiguity Principle - "Students learn better when corresponding words and pictures are presented near rather than far from each other on the page or screen."

- Temporal Contiguity Principle-"Students learn better when corresponding words and pictures are presented simultaneously rather than successively."
- Coherence Principle - "Students learn better when extraneous material is excluded rather than included."
- Individual Differences Principle- "Design effects are stronger for low-knowledge learners than for high knowledge learners, and for high-spatial learners rather than for low-spatial learners."

Challenges to the Application of Principles

Not all research has found that the principles of multimedia learning apply generally outside of laboratory conditions. For example, Muller, Lee, and Sharma found that the coherence principle did not transfer to an authentic learning environment. In their study, adding approximately 50% additional extraneous but interesting material did not result in any significant difference in learner performance.

MULTIMEDIA LITERACY

New literacies generally refers to new forms of literacy made possible by digital technology developments, although new literacies do not necessarily have to involve use of digital technologies to be recognized as such. The term "new literacies" itself is relatively new within the field of literacy studies (the first documented mention of it in an academic article title dates to 1993 in a text by David Buckingham).

Its definition remains open, with new literacies being conceptualized in different ways by different groups of scholars. Accompanying the varying conceptualizations of new literacies, we find a range of terms used by different researchers when referring to new literacies, including *21st century literacies*, *internet literacies*, *digital literacies*, *new media literacies*, *multiliteracies*, *information literacy*, *ICT literacies*, and *computer literacy*. In the *Handbook of New Literacies Research*, Coiro, Knobel, Lankshear, and Leu (2008) note that all these terms "are used to refer to phenomena we would see as falling broadly under a new literacies umbrella" (pg. 10). Commonly recognized examples of new literacies include such practices as instant messaging, blogging, maintaining a website, participating in online social networking spaces, creating and sharing music videos, podcasting and videocasting, photoshopping images and photo sharing, emailing, shopping online, digital storytelling, participating in online discussion lists, emailing and using online chat, conducting and collating online searches, reading, writing and commenting on fan fiction, processing and evaluating online information, creating and sharing digital mashups, etc.

Definitions

The field of new literacies studies is characterized by two theoretically distinct approaches that overlap to some extent. One is informed by cognitive and language processing theories such as cognitive psychology, psycholinguistics, schema theory, metacognition, constructivism, and other similar theories. This orientation includes a particular focus on examining the cognitive and decoding processes involved in comprehending online or digital texts. Donald Leu, a prominent researcher in the field of new literacies, has outlined four defining characteristics of new literacies, according to a largely psycholinguistic orientation. First, new technologies (such as the internet) and the novel literacy tasks that pertain to these new technologies require new skills and strategies to effectively use them. Second, new literacies are a critical component of full participation—civic, economic, and personal—in our increasingly global society. A third component to this approach is new literacies are *deictic*—that is, they change regularly as new technology emerges and older technologies fade away. With this in mind, "what may be important in reading instruction and literacy education is not to teach any single set of new literacies, but rather to teach students how to learn continuously new literacies that will appear during their lifetime." Finally, new literacies are "multiple, multimodal, and multifaceted," and as such, multiple points of view will be most beneficial in attempting to comprehensively analyze them. The second approach to studying new literacies is overtly grounded in a focus on social practices. According to Colin Lankshear and Michele Knobel from this "social practice" perspectives, "new literacies" can refer to "new socially recognized ways of generating, communicating and negotiating meaningful content through the medium of encoded texts within contexts of participation in Discourses (or, as members of Discourses)". From this perspective "new" refers to the presence of two dominant features of contemporary literacy practices.

The first is the use of digital technologies as the means of producing, sharing, accessing and interacting with meaningful content. New literacies typically involve screens and pixels rather than paper and type, and digital code (that renders texts as image, sound, conventional text, and any combination of these within a single process) rather than material print. The second defining feature of new literacies is their highly collaborative, distributed, and participatory nature, as expressions of what Henry Jenkins calls engagement in participatory culture, and Lankshear and Knobel refer to as a distinctive *ethos*.

Research in New Literacy

Research within the field of new literacies is also diverse. A wide range of topics and issues are focused upon, and a broad range of methodologies are used.

Reading and Online Comprehension

One aspect of new literacies that has attracted researchers' attention is school-age children's online reading comprehension. Specifically, researchers are interested in finding the answers to questions such as how reading online differs from traditional print-based reading. In their research, Donald Leu and Julie Coiro attempt to understand how students become adept at online reading, and how students acquire the necessary skills, strategies, and dispositions to comprehend online texts. According to Leu and colleagues, the new literacies of online reading comprehension are based around five defining functions: "These new literacies allow us to use the Internet and other ICT to identify important questions, locate information, critically evaluate the usefulness of that information, synthesize information to answer those questions, and then communicate the answers to others"

Online Fan Fiction and Adolescents

Recent research in the field of new literacies has focused on fan fiction on the internet, especially those stories published online by adolescents. Online fan fiction websites, such as FanFiction.Net, are spaces where fans of all ages, but especially adolescents and younger school-age fans, are able to use these new information and communication technologies (ICTs) to write and craft fictional stories based on their favourite characters in popular media such as movies, television, and graphic novels. Adolescents are participating more and more on these kinds of sites, not only engaging with "pop culture and media, but also with a broad array of literate activities that are aligned with many school-based literacy practices" Of course, adults are just as able to spend time in such online environments—and they do. However, it is interesting to consider that young people were born into a digitally rich world, and thus could be seen as "digital natives," and therefore interact with the online environment in a fundamentally distinct way than an older generation of people, so-called "digital immigrants." As digital natives, adolescents "use the online world to share, evaluate, create, report and programme with each other differently to digital immigrants," and they engage easily and readily with new digital technologies such as instant messaging, file sharing songs and videos, and post all kinds of 'texts,' stories, photos, and videos among them. A central characteristic of digital natives is their "desire to create." Digital natives are engaged in "programming to some extent, whether it be by including a piece of html code that personalizes a MySpace page or creating a Flash animation. They are creating web pages, blogs, avatars and worlds; and, in stark contrast to digital immigrants, digital natives readily report and share their ideas."

Video Games

James Paul Gee described video gaming as a new literacy "in virtue of the ways game design involves a multimodal code comprising images, actions, words, sounds, and movements that players interpret according to gaming conventions". Gee notes that game players participate in their game's world as a form of social practice, especially in real-time strategy games in which players can compete with each other to build on land masses, for instance, and in which they can shape and convey their virtual identities as a certain kind of strategist. Video games continue to use new digital technologies to create the symbols players interpret to encode and decode the meanings that constitute the game. Other researchers have expanded the body of knowledge about video games as a new literacy, particularly as they relate to classroom learning. In one study, researchers examined how video games, supported by discussions and dramatic performances in the classroom, can contribute to the development of narrative thought as demonstrated in written compositions in various contexts. Namely, the researchers engaged primary school children in activities designed to teach them to tell, play out, and write stories based on the most popular video game in their classroom, *Tomb Raider*. Although it was a challenging process at times, researchers discovered that the use of new digital media such as video games actually "complements the use of other written or audiovisual methods [in the classroom] and permits the development of multiple literacies in the classroom."

New Literacies and the Classroom

It has become clear to many researchers in the field that we need to better understand new literacies so that we can answer one critical question: How can educators best incorporate new literacies into the classroom? Kist (2007) observes that in the literature, there are examples of how new literacies can be used in classroom settings—from the use of rap music to anime to digital storytelling, there are already instances of teachers attempting to blend new literacies with traditional literacy practices in the classroom. Kist asks: "Can new literacies indeed 'fit' into how we currently 'do' school?" Kist notes that "the new literacies instruction that does exist often comes only out of the fortitude of lonely pioneers of new literacies." Knobel and Lankshear (2006) argue that if educators and prospective teachers engage in blogging, or participate in "affinity spaces" devoted to practices like fan fiction, video game-playing, music and video remixing, photosharing, and the like, they will better understand how new literacies can better be integrated into worthwhile classroom learning. Leu, Coiro, Castek, Hartman, Henry, & Clemson (2008) have begun to explore the use of a modified instructional model of reciprocal teaching that reflects some of the differences between offline and online reading contexts. In an instructional model known as Internet Reciprocal Teaching, each student has his/her own laptop with access to the Internet and students work in small groups to facilitate interactive group work and discussions about strategy use. In addition, Internet Reciprocal Teaching with online informational resources (as opposed to narrative texts) and strategy instruction on both the common and unique processes by which students navigate through multiple and different texts, rather than the reading of one common text. Teachers and students model their choices about which links are most relevant to a group or individual question through think-alouds. They discuss how to efficiently locate information within different kinds of websites, how to synthesize ideas across multiple texts and media, how to make judgments about the quality of the information and the author's level of expertise, and how to best represent the answers to their questions. Responsibility for monitoring and effectively using these strategies to solve online information problems is gradually released to the students using an instructional scheme with three phases: Phase 1 includes direct, whole class instruction of basic skills and strategies of Internet use; Phase 2 includes group work and the reciprocal exchange of online reading comprehension strategies by students with their peers; and Phase 3 includes online individual inquiry units, sometimes with collaborative efforts involving other students in other classes, perhaps even in other parts of the world, and periodic strategy sessions with groups.

LARGE-SCALE CONCEPT ONTOLOGY FOR MULTIMEDIA

The Large-Scale Concept Ontology for Multimedia (LSCOM) project was a series of workshops held from April 2004 to September 2006 for the purpose of defining a standard formal vocabulary for the annotation and retrieval of video.

LSCOM Mandate

Sponsored by the Disruptive Technology Office (DTO), LSCOM brought together representatives from a variety of research communities, such as multimedia learning, information retrieval, computational linguistics, library science, and knowledge representation, as well as "user" communities such as intelligence agencies and broadcasters, to work collaboratively towards defining a set of 1,000 concepts. Individually, each concept was to meet the following criteria:

- Utility: the concepts must support realistic video retrieval problems
- Feasibility: the concepts are capable or will be capable of detection given the near-term (5 year projected) state of technology
- Observibility: the concepts occur with relatively high frequency in actual video data sets

Jointly, these concepts were to meet the additional criterion of providing broad (domain independent) coverage. High-level target areas for coverage included physical objects, including animate objects (such as people, mobs, and animals), and inanimate objects, ranging from large-scale (such as buildings and highways) to small-scale (such as telephones and applicances); actions and events; locations and settings; and graphics. The effort was led by researchers from Carnegie Mellon University, Columbia University, and IBM.

LSCOM Development Tracks

The project had two main "tracks" — the development and deployment of keyframe annotation tools (performed by CMU and Columbia), and the development of the LSCOM concept hierarchy itself. The second track was executed in two phases: The first consisted in the manual construction of an 884 concept hierarchy, was performed collaboratively among the research and user community representatives. The second track, performed by knowledge representation experts at Cycorp, Inc., involved the mapping of the concepts into the Cyc knowledge base and the use of the Cyc inference engine to semi-automatically refine, correct, and expand the concept hierarchy. The mapping/expansion phase of the project was motivated by a desire to increase breadth—the mapping had the effect of moving from 884 concepts to well past the initial goal of 1000—and to move LSCOM from a one-dimensional hierarchy of concepts, to a full-blown ontology of rich semantic connections.

Project Results

The outputs of the effort include:

1. A "lite" version of the LSCOM concept hierarchy consisting of a subset of 449 concepts.
2. A corpus of 61,901 video keyframes, taken from the 2006 TRECVID data set, annotated using LSCOM "lite."
3. The full LSCOM taxonomy of 2,638 concepts, built semi-automatically by mapping 884 concepts,

manually identified by collaborators, into the Cyc knowledge base, and querying the Cyc inference engine for useful additions.

4. The full LSCOM ontology, in the form of a 2006 ResearchCyc release that contains the LSCOM mappings into the Cyc ontology.

Use of LSCOM in Larger Research Community

Since its release, LSCOM has begun to be used successfully in visual recognition research: Apart from research done by LSCOM project participants, it has been used by independent research in concept extraction from images, and has served as the basis for a video annotation tool.

MULTIMEDIA CITY

Multimedia City (Polish: *Miasteczko Multimedialne*) an innovative project, that has been realized in Nowy S¹cz, in southern Poland. It has started in 2006, on the initiative of leaders and alumnus from WSB-NLU (Wy¿sza Szko³a Biznesu — National Luis University), with Krzysztof Pawlowski and Krzysztof Wnek at the head of. The Multimedia City's main goal is to initiate the cooperation between science and business in testing, incubating and commercialization of innovative projects of new technologies idea to economy.

Project Mission

The aim of the Multimedia City Project is to establish in Nowy S¹cz Center of Innovation working in the field of multimedia and informative system. The strategic goal of Multimedia City is to become one of the most innovative centers in the world, which are working on application of multimedia in education, business and entertainment.

The individual elements of Multimedia City are enabling to implement innovation to economy in accordance with the following stages of innovation chain: fresh ideas and innovative know-how, testing ideas through research and development phase, and implementation and adaptation of the innovative solutions in enterprises.

Current Activities

The following elements of Multimedia City currently operate:

- *Business Incubator MEDIA 3.0*, which within 2 years will invest around 3 mln € in the development of innovative start-ups from the field of IT and multimedia.
- *Animation and 3D Graphics School Drimagine*.
- *The Cluster of Multimedia and Information Systems (MultiCluster)*, that consists of over 60 innovative and highly specialized small and medium-sized companies from the new technologies & new media businesses all over Poland. The Cluster of Multimedia and Information Systems aim is to support and integration of small and medium enterprises sector in the area. Mentioned firms act in multimedia and information

systems market – so called “ creative industry” – and Research and Development Sector.

Numerous companies like Microsoft, CISCO or Lewiatan Business Angels already cooperate with the innovative project from Nowy Sacz.

Realization

The following are wlements of Multimedia City that will operate with the realization of the project:

- *Technology Park* — a complex of buildings (16 thousand square meter) containing laboratories providing the most modern equipment indispensable to do research and developmental works in the field of multimedia and IT.
- *Research and Development Centre* — will carry a range of activities focused on searching for multimedia new applications. In the centre the ideas created, among other things, in the Cluster of Multimedia and Information Systems will be developed. Infrastructure and equipment will enable research and testing of any multimedia or IT project in the R&D Centre laboratories.
- *Business Incubator* — will enable the new, innovative ideas in the field of multimedia to be expanded. Furthermore, the ideas will be developed coming from different sources such as Research and Development Centre, firms, students and trainees who carry out

studies using Multimedia City infrastructure and facilities.
- *Multicluster.*

Multimedia Technology Support

The following are areas of multimedia technologies that the project strongly supports:

- Mobile Technologies such as creation of localization service system and authorization systems based on the integrated technology of internet services.
- Computer animation 3D used in computer games, special effects in films, visualization of data and simulations.
- Computer games used in education, business and entertainment.
- Internet, development of technologies such as Web 2.0, CSS, PHP, Java to be used in different purposes and fields.
- E-marketing, e-advertising, use of multimedia in marketing and advertisement.
- E-learning, increasing multimedia technology usage in education.

Technology Park

On 24 August 2010, Multimedia City signed an agreement with the Polish Agency for Enterprise Development for a bailout to build the Technology Park. Multimedia City will be given 25 mln € for its infrastructure. Thanks to the European funds from the Innovative Economy Operation Programme (indicative list) the modern research and development infrastructure of the Technology park will be built till 2012. That investment will provide a chance for Nowy S¹cz to become the Polish Innovation Valley.

In September 2010, building of a complex infrastructure for the Technology Park (16 thousand square meter) has started.

In Multimedia City's main building will be found specialist technology laboratories of amongst others: 3D and special effects, virtual reality, post-production, sound, motion capture as well as modern offices surface.

MULTIMEDIA AUSTRALIA

Multimedia Australia is a software publishing and multimedia company based in Brisbane, Australia that services over 90 countries. The company was founded in May 2001.

The company is primarily known for its role in establishing the Software Industry Professionals association, its line of software used in Web development, and its large network of websites. Multimedia Australia also became one of the first Australian IT companies to generate its own electricity using PV solar panels for powering some of its computer equipment in September 2007 after conducting a trial earlier in the same year.

Corporate Social Responsibility

Multimedia Australia's corporate social responsibility initiatives have included supporting disaster relief organizations, providing free software to schools and educational institutions (which has since been changed with more products provided at academic rates) and supporting various causes such as Earth Hour.

History

Upon its founding in May 2001 Multimedia Australia started as a publisher of Web design and development software. Its first product was the BestAddress HTML Editor which is still on the market today and which has been recognised with over 40 awards and had several reviews written about it. Multimedia Australia's news archives show that it continued operating solely as a software publishing company until February 2004, when it launched the first of its information-based websites which would eventually form part of the large website network that the company operates today. In June 2005 Multimedia Australia acquired Hixus Software. Later in October that same year the company began offering Web hosting and search engine submission services. In May 2006 Multimedia Australia founded the Software Industry Professionals group as one of its corporate social responsibility initiatives with the aim of enabling people in the software industry to network with each other and access tools and resources to help them be more successful in the industry. Software Industry Professionals has subsequently become one of the largest membership associations in the software industry with over 2000 members in 93 countries as of November 2008. Multimedia Australia still supports Software Industry Professionals today. In July 2008 the company entered the mobile device content publishing market. Multimedia Australia expanded into Web design in January 2009. In March 2010, Multimedia Australia launched the Australia Times newspaper site. In an Australian first, on May 5, 2010, the Australia Times become the first Australian newspaper to release newspaper reading software with its Australia Times Reader download.

Key Software Products

- Acclaim CMS
- Aurora Web Editor
- BestAddress HTML Editor
- DigitalAccess FTP
- PageLock Website Copy Protection
- Web Palette Pro
- Privacy Solver (No longer sold)
- Clipsafe Clipboard Backup (No longer sold)

In total the company sells 16 software products, either under the *Multimedia Australia* brand name, its other divisions, or products that it obtained through the acquisition of another business.

Key Services

- Web design
- Web hosting
- Domain name registrations
- Search Engine Optimisation
- Software promotion
- Website network

MULTIMEDIA - PLANNING AND MANAGING FOR A BETTER PRODUCT

Planning

A multimedia product, such as Power Point, can really get the students going in a computer lab. Presentations can become extremely complex and help is needed often when students are in the middle of creating special effects. There really is not much to be accomplished if there has been no planning between teacher and computer tech. Students have to have research geared towards this product; it will not work if they bring in a typed report intending to produce a multimedia product. Use of graphic organizers are key to the success of a multimedia project, because students must become aware that they are NOT writing a series of paragraphs. They will be presenting a series of main ideas, facts, or short descriptions. The presentation will be a series of charts, and student must have a good idea of how their topic is to be broken up. Bibliographies should also be complete before beginning the project. Students should have all research and graphic organizers complete before beginning Power Point, or any other multimedia software, because they will not have a good plan for the total presentation. Teachers can review the drafts of the text in hand written form or typed, printed format. Power Point allows the easy printing of text only for teacher review.

Types of Multimedia Projects

Teachers should decide which type of presentation will apply, or if the students will make this choice. The three types are:

- Self Playing presentations. With this type, slides advance automatically and all special effects play automatically. There is no manual intervention at all, and the presentation becomes a "show."
- Manually advanced presentations with linear progression. This type of presentation will advance slides and text only on the click of the mouse. Special effects can be set to go off automatically or by a mouse click. Students required to give an oral presentation accompanying their multimedia presentation should use this type. Example of topics with linear progression are biographies or History.
- Manually advanced presentation without linear progression, or "interactive." This type of presentation is most difficult for students to plan in advance if they haven't seen it done already. The front slide will have hyper links which can be clicked in any order by the

presenter. When clicked, another slide will appear which has specific information, special effects, and a place to click which takes the presentation back to the front slide or on to another. Students should be required to make a "map" in advance, with lines connecting all slides (boxes) which are to be connected by action buttons. Otherwise it becomes difficult to visualize. This type of presentation are for topics where there is no particular order that information should be presented. Examples are parts of a plant, math projects, etc. This type of presentation can be created with an accompanying oral talk in mind, or it can be created as a presentation which students will take turns sitting down and studying it interactively. This is a particularly good format for students [or teachers] who want to create interactive quizzes. Power Point has a feature for action buttons which when clicked, will take the viewer back to whatever slide brought the viewer to that point. This means that a generic "wrong answer" slide can be prepared which will work again and again to return the viewer to whichever slide they came from. Buttons can also activate sound effects, or other actions.

Graphics

Teachers should have an idea ahead of time what graphics, if any, the students are to use for the project. Many choices are:

- Maps, photos or other graphics can be found and downloaded from the Internet during the class computer lab time. Students will be taught how to search for, download and insert graphics into their documents. Permission is not necessary as long as the use of the graphic remains within the confines of classroom instruction.
- Students can be taught how to use the drawing tools to alter graphics (e.g., to put an arrow mark on a map graphic), or to create simple graphics themselves using Word Art, autoshapes, etc.
- Students can use scanners to scan in graphics found in books or art work they themselves have created. Students will be taught how to scan, save, manipulate and insert the graphic into their documents.
- Students can create their own graphics using Windows paint or some other software.
- Students can use a digital camera to take a photo. Students will be taught how to download the image from the camera, save it, manipulate it and insert it into their documents.
- Students can use conventional clipart programmes or encyclopedia CDs to find suitable graphics. Students will be taught how to search clipart or encyclopedia CDs, save the image and insert it into their documents.

Pulling Multimedia in Technology Skills

A multimedia presentation has a number of ways that the media can be manipulated to support or enhance the topic. Students should be taught the various ways and encouraged to create presentations which are a "total experience." The available facets of multimedia are discussed below:

- Graphical design and the use of colour has a powerful effect in a multimedia presentation. Students should be taught (if time allows) how colours are associated with moods. Teachers should decide in advance how much time to allocate to this topic. Areas of the art curriculum can be emphasized here which deal with use of imagery, lines and colour.
- Background music also has a powerful effect to enhance and support a topic. Teachers should think about and establish rules for use of music and how much freedom student preference should be given. Many music files are available on the internet which allows students to conform to the period/ethnic theme of the presentation. If time allows, students should be taught how to coordinate music patterns to animation patterns. Music links can be found here.
- Sound effects are available on the Internet and some in Power Point. Used sparingly and with careful selection of quality, they add a tremendous depth of feeling to a presentation. Due to the inclusion of

inappropriate sounds in all sound web sites I've seen, I do not have a link of sound effects for students to browse through as I do for music. However, teachers can search on "sound effects" and come up with a number of sites having a rich variety of sound effects most of which are designated as open domain files.

- Students can also use a microphone to create their own sound effects. Unexpectedly hearing student voices during a presentation jars the "viewer" into attention with stunning exactness. Again, teachers should decide in advance how much time to allocate to this activity. Due to the need for complete silence, the recording session usually has to be the same period for everybody in the class. The students take turns using the microphones. I have also had a "free period" designated for sound recording when a whole class needed to do this.
- Video also adds fascination to a presentation. Video clips can be downloaded or student created video can be added. If students create their own clips, they must be encouraged to keep the time to 5 seconds or below as the file sizes can become enormous.
- Buttons add a fascination for the viewer. If free access to the total presentation is granted by generous application of buttons, the viewer becomes engaged

in a learning experience. Students should be encouraged to explore interactive presentations.

- Animation and timing. It is critical that students be taught how important it is that the motion and progression of the multimedia be controled in a manner that allows the viewer to digest the content. This means readable text, supporting graphics, animation that *directs*, not *distracts*.

Many technology skills can be chosen selectively for introduction during a multimedia project. These include:

- Text manipulation (font, size, bold, italic, underline, justification, colour, text art, text boxes, spell checking, grammar checking, copying, cutting, pasting, deleting). Text needs to be readable, inviting (ie, NOT long paragraphs) and attractive in multimedia.
- Methods to save and retrieve files (the correct server, path, file name, file extension, folder, drive)
- Object formatting (text boxes, sound, video, and graphics boxes, formatted for animation, autoplay, manual advance, autostart, loop, timing, dim/flash)
- Graphics manipulation (inserting, resizing, cropping, framing, copying, cutting, pasting, deleting, moving, grouping)
- Use of peripherals (microphone, CD ROM, scanner, digital camera, and printer)

- Other Technology Skills: tool bars, short cut icons, menu system, help system, short cut keys, task switching, and other features which can be utilized during word processing. Certainly if there is much searching the Internet for graphics, video or music, task switching can become second nature to the students.