Lecture 2: Access Networks

Dia 3

The access network forms in important part of an end-to-end communication network. Although the boundaries of an access network are not always very clear, in general one may consider that part of the overall network that is connecting a customer terminal or customer premises network with the core network as an access network. In general the access network is not incorporating any routing or switching functionality.
The figure illustrates a generic access network connected to a generic core network via a gateway. The core network will support the exchange of information between different gateways.

Dia 4

Different types of access networks may be distinguished based on the transmission medium used and on the services offered. In the past there was a clear separation between different types of access networks :
The twisted pair access was used for access to the telephone network (PSTN or Public Switched Telephone Network). The terminal is a telephone and the gateway is a telephone switch : local exchange (LEX).
The coaxial cable was used for access to the CATV network (Community Antenna TeleVision). The terminal is a television and the gateway is a head end (where the TV-programs are captured and multiplexed on the coaxial cable).
An air interface is used for access to a mobile network. The gateway is the Mobile Switching Center (MSC) giving access to the PSTN network.
In case of computer networks getting access to the Internet, the situation is less clear. One could consider the access network as the LAN or Local Area Network that is connected to an edge router (as shown in the figure). One could also consider the connection from the edge router or from a home computer to the public Internet (via cable modem, ADSL modem, ISDN modem, …). In the latter case, the physical infrastructures are the access networks mentioned above (twisted pair, coax, air).

Dia 6

There are three important guided transmission media used for transmitting information : twisted pair (used to connect your home to the telephone network), coaxial cable (used to connect your home with the TV-distribution network : CATV) and optical fiber (used to interconnect large telephone exchanges). An unguided transmission medium is the air, used for wireless access (not discussed here).
A twisted pair consists of two isolated copper wires which are twisted together. By twisting them, one will be less sensitive to external noise sources. A number of twisted pairs will be bundled in a cable.
A coaxial cable consists of one central copper wire surrounded with a copper shield. Between the copper wire and the shield one is using an insulator. In this way one is able to have a very good shielding against external noise sources and at the same time the bandwidth of this medium is increased.
An optical fiber consists of a glass wire with in the center a higher refractive index. In this way the light will be guided in the fiber (by the principle of total internal reflection).

Dia 7

A signal transmitted over a transmission medium (twisted pair, coax, fiber) will be attenuated. This attenuation is commonly expressed in dB/km. A value of 3 dB/km means that the power of the signal is reduced by a factor 2 when it is transmitted over a distance of 1 km over the specific transmission medium. When transmitting over 3 km one will observe an attenuation by a factor 8. This attenuation strongly depends on the frequency used, as indicated in the figure. We observe that the attenuation for twisted pair is much higher than for coaxial cable. Optical fiber has the lowest attenuation for frequencies (or wavelengths) in the infrared range. As a result, optical fiber has the largest capacity to transmit signals. Today total bitrates of >5 Tbit/s (= 5. 1012 bit/s) on a single fiber have been demonstrated.
When a signal is attenuated, one has to amplify it again in order to be able to bridge a certain distance. The distance of these repeaters is indicated in the table.
Note : a distinction is made between a loaded and an unloaded twisted pair. The loaded twisted pair is using some passive elements (coil) in order to obtain a fairly flat frequency response in the frequency range for analog voice signals (0.3 to 3.4 kHz). For higher frequencies a large attenuation is introduced (making this useless for higher frequencies). An “unloaded” twisted pair will have a larger bandwidth (but not a flat response). Note that the twisted pair in both cases is the same, but the use of passive elements is different (which could be easily removed).

Dia 8

As was shown in the previous slides, a transmission medium is limited in bandwidth. This figure illustrates the deformation of pulses when we transmit them over a bandwidth limited system. The lower the available bandwidth, the worse the approximation of the pulses.
From this it becomes clear that the maximum bitrate one can transmit over a twisted pair will be much lower than on a coax cable. The optical fiber allows the largest bitrate.

Dia 9

The table shows some characteristics of different types of twisted pairs : UTP (Unshielded Twisted Pair : ordinary telephone wires) and STP (Shielded Twisted Pair : a metallic braid or sheathing is reducing the interference noise). Two important categories of UTP are considered : category 3 (twist distance : 7.5 … 10 cm) and category 5 (twist distance : 0.6 … 0.85 cm). The Shielded Twisted Pair is clearly performing better.
The first parameter in the table is the attenuation (dB /100 m). The second parameter is the near-end crosstalk. This is the amount of a signal transmitted on one twisted pair that will be coupled to another twisted pair (both in the same cable). The goal is to have a very low crosstalk. A large value (in dB) means low crosstalk. If the crosstalk between two twisted pairs is 30 dB, this means that a signal with strength 1 transmitted on one twisted pair will result in an unwanted signal of 0.001 on the other twisted pair. A higher value (e.g. 50 dB) results in a better crosstalk behaviour (only 0.00001 is carried on the other twisted pair).

Dia 10

In terms of optical fibers, one observes 3 important types : step-index multimode, graded index multimode and single mode (see figure). Without going into the details, one can say that the single (or monomode) fiber is the best and it is in general used in telecommunication networks. Only for short distances one is using multimode fiber (the graded index type).

Dia 11

In this figure we observe a number of important frequency bands for electromagnetic waves. Note that electromagnetic waves are e.g. used to transmit the signal from a GSM to the antenna, or they are used in a microwave furnace. Light is also an electromagnetic wave (but with a very high frequency).
A very important property of electromagnetic waves is that they propagate with the speed of light. The value of the speed of light (in vacuum or air) is $c = 3.10^8$ m/s. Note that this speed is different in e.g. optical fiber (appr. 2.10^8 m/s). If we look at a propagating wave, we can observe a wavelength λ which is the period (in space) of the signal. We have a relationship between frequency and wavelength : frequency x wavelength = velocity (or $f \times \lambda = c$).

Examples :
Red light with a wavelength of 600 nm (or 0.6 mm) gives a frequency of 5.10^{14} Hz (=500 THz*).
A GSM works at 900 MHz (or 1.8 GHz) and is using a wavelength of 33.3 cm (or 16.6 cm).
An FM radio works around 100 MHz and is using a wavelength of 3 m.
* kilo or $k=10^3$, mega or $M=10^6$, giga or $G=10^9$, tera or $T=10^{12}$, peta or $P=10^{15}$

Dia 12

Transmitting a signal over a transmission channel (coaxial cable, twisted pair), introduces noise. This will be added to the input signal, resulting in (sometimes) severe distortion of the signal. As a result, some wrong bits may be received. This results in a Bit Error Rate (BER), which is the average number of wrong bits received (averaged over a long observation time). With optical fiber transmission one typically gets BER of 10-9 or better which means that one receives on average 1 wrong bit on a total of 109 received bits.
The noise introduced in a transmission channel will result in a limitation on the number of bits that can be transmitted over the channel. This is expressed by the Shannon capacity limit (see next slide).

Dia 13

From the Shannon capacity limit we learn about the maximum bitrate C (bit/s) one can transport over a bandwidth limited channel with noise*.
C = B log2(1+SNR)
The bandwidth B is expressed in Hz. The noise is specified as a relative measure with respect to the signal strength : SNR or Signal to Noise Ratio = signal power/noise power. In many cases this is expressed in decibel (dB) : SNRdB = 10 log10 (signal power/noise power). When the signal power is 1 and the noise intensity is 0.001, we have a SNR ratio of $1000=10^3$ or expressed in dB we have a SNRdB of 30 dB (this is a typical value for telephone connections, see also in a previous slide).
If we take an example of B=4 kHz and SNRdB=30 dB, we obtain a maximum capacity of C = 40 kbit/s.

*Note that for a channel without noise the bitrate is in principle infinite : “you can send every second a signal with an amplitude specified with an infinite accuracy”. For example at t=0 you send a sample with amplitude 1.384932829048321… (infinite number) and at t=1 you send a next sample with amplitude 2.09473829103858392… and so on. This gives indeed an infinite transfer of information because you suppose that you can receive this signal with infinite accuracy (no noise is added). If there is a noise signal with amplitude e.g. 0.001, then the digits after 0.001 do not have any value any more : due to the noise you do not know if the original signal was 1.384 or 1.385.

Dia 15

The telephone access network makes use of twisted pair to transmit the voice signal (0.3 - 3.4 kHz bandwidth) from the LEX to the user. The topology is a star with a typical radius between 1 and 10 km. The transmission is bidirectional point-to-point. In a typical access network one had over 10000 users connected to the same LEX and the penetration (number of users connected) is larger than 90%.
It is important to remember that about 80% of the network transmission cost is in the access network.

Dia 16

The telephone access network makes use of cables with twisted pairs (SDM or Space Division Multiplexing) in a physical tree and branch structure (the twisted pairs form a logical star topology).
Very big cables leave the Local Exchange (LEX) and are further split in smaller cables. In the distribution points (or street cabinets) the cables are finally connected to the homes (see next page).

Dia 17

The physical roll-out of the part between the distribution point and the subscriber premises is not a star. For economical reasons a number of twisted pairs are put in one cable and laid along a row of houses.
Each household is provided with two twisted pairs, which of course results in the possibility of having two connections to the telephony network. When there is demand for a second connection there is no need to dig up (expensive) the streets again.
When a subscriber has to be connected to the network the local telephone operator performs a connection between a free twisted pair in the feeder cable and the twisted pair which is laid from street cabinet to the subscribers wall plug. This connection is accomplished by applying a jumper in the street cabinet.

Dia 18

Some examples are given of the twisted pair cables used in the exchange buildings (top picture). Distribution frames (where the twisted pairs are terminated) are shown at the bottom.

Dia 19

As more multimedia services become available today, the use of voice-band modems is no longer sufficient to transport the “bandwidth-hungry” applications such as Internet, file downloading, email, digital TV, interactive games etc.

More capacity and higher line speeds need to be offered by the operator. As the installed base represents a huge investment, a lot of R&D was spent the last decade to develop techniques and scenario's for upgrading the existing networks.
The digitalisation of the transmission was already introduced with the ISDN service (Integrated Services Digital Network). At the same time development of better, more efficient modulation techniques allows a better usage of the available transmission bandwidth of the TP medium (e.g. QAM or Quadrature Amplitude Modulation). The frequency bandwidth itself, used for the transmission of information on the TP, has also been extended up to the MHz regions (removal of passive elements). As losses are higher on the TP in this frequency range, the distances that can be covered without amplification become shorter. This problem is being alleviated by the operator by "upgrading" the TP network itself by gradually introducing optical fibre in the network. It is not sufficient to increase the capacity of the twisted pair access network, it is also important to connect it to a "multimedia network". Since no real multimedia network exists today, additional services are offered by an interconnection to the Internet (mainly for data services).
Note that while introducing more multimedia services on the TP network, using several advanced techniques and evolution scenarios, the classical telephone service still has to be offered to the customers.

Dia 20

Introduction of optical fibre as transmission medium is done gradually. The "large capacity" optical fibres are introduced up to a certain distance in the access network (keeping twisted pair for the "last mile" to the customer). At that point, conversion from optical to electrical information signals needs to be supported. This network element is called an Optical Network Unit (ONU).
The "last mile" distance to the customer is still covered by the TP : this is the major part of the installed network and it would be very expensive to replace it with fiber. Depending on the bandwidth and capacity offered, this distance can vary from a few km down to a few hundred of meters.

Dia 21

Introduction of optical fibre as transmission medium is done gradually. The "large capacity" optical fibres are introduced up to a certain distance in the access network (keeping twisted pair for the "last mile" to the customer). At that point, conversion from optical to electrical information signals needs to be supported. This network element is called an Optical Network Unit (ONU).
The "last mile" distance to the customer is still covered by the TP : this is the major part of the installed network and it would be very expensive to replace it with fiber. Depending on the bandwidth and capacity offered, this distance can vary from a few km down to a few hundred of meters.

Dia 22

This figure illustrates the bitrate of a number of modems.
A first category of modems uses the voice-band frequency range (from 0.3 to 3.4 kHz). Today bitrates up to 56.6 kbit/s are possible (Close to Shannon limit).
A second category uses a much higher bandwidth (up to a few MHz) and results in much higher bitrates. They belong to the xDSL line systems (DSL : Digital Subscriber Line).

Dia 23

A number of broadband modems is illustrated in the table. The following characteristics are indicated : bitrate, modulation format, required bandwidth, number of twisted pairs used, mode of operation. It is also indicated whether the technology may work as an overlay on the analog PSTN or digital ISDN access. It is observed that only ADSL and VDSL may form an overlay (they may be operated in parallel with PSTN or ISDN on the same twisted pair).
The first technology (T1/E1) is supporting a primary multiplexer signal, well known from classical telephony multiplexing. The bitrate used in Europe is 2 Mbit/s.
An ISDN (Integrated Service Digital Network) line is running in full duplex mode on a single twisted pair. EC means that Echo Canceling is used.
The remaining technologies are xDSL or Digital Subscriber Line technologies.
HDSL (High Speed DSL) is offering access to the network similar to a T1/E1 (note however the difference in modulation, required frequency band and number of twisted pairs). The lower bandwidth is obtained by splitting the 2 Mbit/s signal over 2 (or 3) full duplex twisted pairs.
ADSL (Asymmetric DSL) is providing an asymmetric access to the network. This was specially designed for residential Internet access which is in nature asymmetric : sending HTTP requests for web access (=upstream from user to network) does not require a lot of bandwidth, compared to the download of the web site (downstream from network to user). The available bitrate (and associated required bandwidth) will depend on the distance to be bridged.
VDSL (Very high speed DSL) further increases the bitrate but over a shorter distance.
HDSL2 and SHDSL are further developments of HDSL.

Dia 24

The figure illustrates the downstream bitrates versus the distance for ADSL and short versus long range VDSL. The trend is clear : the higher the bitrate, the smaller the distance that can be bridged. Note : 1 meter = 3 feet.

Dia 25

The figure shows the percentage of households that have broadband access (DSL, cable, …). The region considered comprises 387 million people and 163 million households.
Note that the difference between Belgium, The Netherlands on one hand and Germany and Italy on the other hand, correspond to a time difference of about 1,5 years. This is also illustrated in the growth figures below. The differences will gradually disappear over the coming years.
Tables from "Heavy Reading, Vol. 3, No. 5, March 2005 : Next-Generation Broadband in Europe"

Dia 26

The table shows some recent trends in next generation DSL roll-outs.
Table from "Heavy Reading, Vol. 3, No. 5, March 2005 : Next-Generation Broadband in Europe"

Dia 27

The length of the local loop (length of twisted pair) varies largely between countries. This results in different opportunities.
Tables from “Heavy Reading, Vol. 3, No. 5, March 2005 : Next-Generation Broadband in Europe”.
Example: ADSL2+ in Europe: 56% of the homes is within the 2 km range, for the READSL 97 % is within 6 km.

Dia 28

The figure shows a (risky) extrapolation of recent trends in access speed. The blue line illustrates the actual connection rate and the brown line indicates the straight line extrapolation. This corresponds with a bandwidth increase by a factor of 2.3 every 2 years. Of course one has to be careful with extrapolation of these data. Table from “Heavy Reading, Vol. 3, No. 5, March 2005 : Next-Generation Broadband in Europe”
The figure below gives a similar view, but somewhat more conservative.

Dia 30

The figure shows the typical architecture of an ADSL based access network. At home a computer or multiple computers are connected (via a switch or hub) towards an ADSL modem (typically via an Ethernet interface or USB port). A POTS splitter will is used to separate the telephone (POTS) signal and the ADSL signal. A Twisted Pair (TP) will connect the home (Remote side) to the access multiplexer (Central side). The twisted pair is terminated in the DSLAM (DSL Access Multiplexer). The DSLAM will connect to the access aggregation network using a L2 technology (today mostly ATM or Asynchronous Transfer Mode, but more recently also using Ethernet). The L2 network is connected to a BRAS (Broadband Remote Access Server) where e.g. PPP connections are terminated. The BRAS will connect to the public Internet or to other service providers.
Details of the different blocks will be discussed in more detail later in this chapter.

Dia 31

The system reference model illustrates the functional blocks required to provide ADSL services. It is part of the ITU-T (International Telecommunication Union – Telecommunication standardization sector) recommendation G.992.1 : Asymmetric digital subscriber line (ADSL) transceivers (dated 06/1999). Parts of the description of ADSL are based on this recommendation.
The top figure gives the basic functional blocks and the standard interfaces between these functional blocks (this approach is very common in access networks, e.g. ISDN).
On the bottom figure from right to left we observe four rectangles:
1. Home environment with SM (Service Module. e.g. PC), customers premises network (e.g. Ethernet or IEEE 802.11), and part of the modem functionality (physical layer towards home network and limited switching capacity (e.g. ATM related). This includes the NT2, TA and User Terminal functional blocks and the T/S interface.
2. The NT1 block (Network Termination 1) includes the ATU-R (ADSL Transceiver Unit – Remote) and the POTS splitter R with a h-p (high-pass) and a l-p (low-pass) filter. We observe some other interfaces (T-R, U-R2 and U-R).

3. At the DSLAM site we observe similar functions (POTS splitter, ATU-C or ATU Central office). Note that at the DSLAM side a large number of end users will be connected and multiplexed towards the broadband network.
4. The IP traffic will be connected to the broadband network and the classical telephony will be connected to the narrowband network.

Dia 33

The overall bitrate over the DSL line is asymmetric (this is why one talks about Asymmetric DSL). Downlink is typically in the range 1 to 6.784 Mbit/s and uplink typically in the range of 64 to 640 kbit/s. This overall capacity is split in several bearer channels, as indicated in the figure. We observe 4 downlink simplex bearer channels (AS0, AS1, AS2 and AS3) and 3 duplex bearer channels (LS0, LS1 and LS2). In addition there is room for operation, maintenance and control. There is a minimum requirement of having at least the AS0 and LS0 channel.
The bitrates of these channels should be a multiple of 32 kbit/s. The table illustrates the ranges the different channels may occupy (if they are present). The AS0 channel may occupy a range of 32 kbit/s (1 x 32 kbit/s) up to 6144 kbit/s (192 x 32 kbit/s).

Dia 34

The figure shows a block diagram of an ATU-C **transmitter** (ATU-C : ADSL Transceiver Unit - Central office) for the downstream transport of ATM data*. The figure illustrates the functional blocks and interfaces. As can be observed, a large part is physical layer related.
Two paths are shown between the Mux/Sync control and Tone ordering; the "fast" path provides low latency; the interleaved path provides very low error rate and greater latency. An ADSL system supporting ATM transport shall be capable of operating in a single latency mode, in which all user data is allocated to one path (i.e. fast or interleaved). An ADSL system supporting ATM transport may be capable of operating in an optional dual latency mode, in which user data is allocated to both paths (i.e. fast and interleaved). ADSL equipment shall support at least bearer channel AS0, support of AS1 is optional.
<u>Some abbreviations:</u>

NTR: Network Timing Reference (8 kHz reference to be transmitted downstream)
OAM: Operation Administration and Maintenance
EOC: Embedded Operations Channel
AOC: ADSL Overhead Control Channel
FEC: Forward Error Correction
Cell TC : Cell Transmission Convergence
CRC: Cyclic Redundancy Check
IDFT: Inverse Discrete Fourier Transform (use of 512 coefficients)
DAC: Digital Analog Conversion
* ADSL is using ATM (Asynchronous Transfer Mode) cells to transport information. This are very short packets (53 bytes) with 5 bytes overhead and 48 bytes payload.
Note: at the ATU-R, a similar functionality is observed but in receiver mode.

Dia 35

This figure shows the return channels for an ATM based ADSL system (at the ATU-R or ADSL Transceiver Unit – Remote). In this case LS0 is always required, LS1 is optional. The LSx channels are used in simplex mode (although they are duplex capable).
Note: at the ATU-C, a similar functionality is observed but in receiver mode.

Dia 38

The ADSL PHY was designed so that it could peacefully co-exist with the standard POTS spectrum. The two services can co-exist because the ADSL spectrum only uses the frequencies above POTS. The POTS spectrum goes from near DC to approximately 4 kHz. A frequency guard band is placed between the POTS spectrum and the ADSL spectrum to help avoid interference. The ADSL spectrum starts above the POTS band and extends up to approximately 1.1 MHz. The lower part of the ADSL spectrum is for upstream transmission (from the customer to the Central Office) and the upper part of the spectrum is for downstream transmission. There are actually two different ways that the upstream and downstream spectra can be arranged:
In a frequency division multiplexed (FDM) system (or FDD or Frequency Division Duplex), the upstream and downstream spectra use separate frequency ranges. They can vary for different implementations, but typically the upstream band is from 26 to 138 kHz and the downstream band is from 138 kHz to 1104 kHz. This system is free from the occurrence of a type of interference called self-crosstalk (the figures also illustrate the power spectral density (PSD) allowed in the two directions). One drawback, however, is that the downstream bandwidth is reduced in comparison to an echo-cancelled system.
An echo-cancelled system allows the downstream band to overlap with the upstream band. The upstream band still uses the same frequencies, but the downstream band can now extend over the upstream band. The main advantage of this system is that it significantly extends the available downstream bandwidth. However, it does require echo-canceling circuitry due to the full-duplex transmission. In addition, the presence of self-crosstalk causes additional interference.

Dia 39

ADSL is using the DMT technique (Discrete Multi Tone). The DMT technique employs frequency division multiplexing by dividing the bandwidth of a twisted pair is smaller bands (each with a central carrier). The carriers are spaced at 4.3125 KHz intervals, with 224 of those used for the downstream data in the range from 138 KHz to 1.1 MHz and 25 of those used for the upstream data in the range from 26 KHz to 133 KHz (in case of FDD).
The figure indicates the DMT spectrum with indication of the POTS band, upstream pilot tone, downstream pilot tone, subcarrier spacing, and number of subcarriers for the upstream and downstream direction. Dividing the available bandwidth into a set of independent, orthogonal subchannels are the key to DMT performance. By measuring the SNR (Signal to Noise Ratio) of each subchannel and then assigning a number of bits based on its quality, DMT transmits data on subcarriers with good SNRs and avoids regions of the frequency spectrum that are too noisy or severely attenuated. The underlying modulation technique is based on quadrature amplitude modulation (QAM). Each subchannel is 4.3125 kHz wide and is capable of carrying up to 15 bits.

The stream of data bits is first divided into several parallel bit streams, and these are used to modulate a sub-carrier each, independently of the others. The bandwidth efficiency attainable in each of these sub bands extends from 0 to 15 b/s/ Hz (depending on the SNR of the subchannel). At low frequencies, where copper wire attenuation is low and signal-to-noise (SNR) ratio is high, it is common to use a very dense constellation supporting 10 b/s/ Hz or more. In unfavourable line conditions, the modulation can be relaxed to accommodate a lower SNR, typically 4 b/s/ Hz or less.
The advantage of this technique is that it adapts to the line condition in the many narrow sub bands by adjusting the constellation density accordingly, and thus avoids wasting energy in those band.

Dia 43

QAM is a technique using a complex hybrid of phase (or ‘quadrature’) as well as amplitude modulation. The figure shows a simple eight-state form of QAM in which each line signal state represents a 3-bit signal. The eight signal states are a combination of four different relative phases and two different amplitude levels. The table relates the individual 3-bit patterns to the particular phases and amplitudes of the signals that represent them. Note the top right figure illustrating the actual line signal pattern that would result if we would send the signals in the table consecutively as shown.

Dia 44

The QAM modulation is using amplitude and phase to make a difference between the different symbols. This can be represented in the complex plane where each point is represented by a complex number (Z = x + jy).
An example for 16-QAM is shown in the table and the figure.

Dia 45

To have a bigger constellation size it is necessary to have a better SNR in order to hold a certain bit error rate. ADSL can go up to 16 QAM modulation (but this will be only possible in case of a very good SNR).
The examples show general noise impact and the impact of phase noise.

Dia 46

ADSL is using a baud rate of 4000 symbols per second (per carrier) or every 0.250 msec a symbol is transmitted on a specific carrier.
Because there are maximally 256 carriers, a maximum of 256 symbols may be transmitted at the same time (in practice this is of course lower because the lowest carriers are used for POTS and because in the FDD the upstream and downstream carriers are different).

Dia 47

The implementation is done using the Inverse Discrete Fourier Transform (IDTF). This will transform the frequency samples (the complex numbers) into (real) time samples. This requires however the use of and extended Z vector (extended with the conjugate values in order to obtain real time values xk). The IDFT will generate (max) 512 time samples in parallel that will be converted in a serial stream. This stream will be converted into an analogue signal that is transmitted over the twisted pair. An example is given in the next slide.

Implementation is typically done using the FFT (Fast Fourier Transform) algorithm.

Dia 48

An example of the frequency domain and time domain is shown (for 128 constellation points).

Dia 49

It is important to know for each of the channels the number of bits that can be transported (will depend on the SNR in that channel, see also Shannon). In order to know the constellations (and the gain factor one should use), the down-stream channel is tested by a broadband pseudo random signal during initialization. The ATU-R receiver calculates the maximum number of bits per symbol that each down-stream channel can support. The target error rate is 10-7 and the performance margin is 6 dB. A table is sent back to the ATU-C receiver with bit allocation, bk, and gain adjustment factor, gk. { bk, gk; k=1, 255 }. When both bk and gk are zero, carrier # k is not used permanently. When b is zero and g is unity, carrier # k is not used temporarily. Gross gain adjustment of 6 dB may be required for carriers above carrier #51. Fine gain adjustment of 1.5 dB may be required to equalize the expected error rate performance across the tones

Dia 52

ADSL uses the superframe structure shown in the figure. Each superframe is composed of 68 data frames, numbered from 0 to 67, which are encoded and modulated into DMT symbols. They are followed by a synchronization symbol, which carries no user or overhead bit-level data and is inserted by the modulator to establish superframe boundaries.
From the bit-level and user data perspective, the DMT symbol rate is 4000 baud (period = 250 µs), but in order to allow for the insertion of the synchronization symbol the transmitted DMT symbol rate is 69/68 × 4000 baud.
Each data frame within the superframe contains data from the fast buffer and the interleaved buffer. The size of each buffer depends on the assignment of bearer channels made during initialization.

Dia 53

Because of the addition of FEC redundancy bytes and data interleaving, the data frames (i.e. bit-level data prior to constellation encoding) have different structural appearance at the three reference points through the transmitter (see below). The reference points for which data framing is shown, are:
• A (Mux data frame): the multiplexed, synchronized data after the CRC has been inserted. Mux data frames shall be generated at a nominal rate of 4 kbaud (i.e. every 0.250 µs).
• B (FEC output data frame): the data frame generated at the output of the FEC encoder at the DMT symbol rate, where an FEC block may span more than one DMT symbol period (in case of the interleaved data buffer).
• C (constellation encoder input data frame): the data frame presented to the constellation coder.

Dia 54

The figure illustrates the frame structure of the fast data buffer at reference point A (without FEC) and at reference point B/C (with FEC added at the end).
Fast Byte : The synchronization byte of the fast data buffer (indicated as "fast byte") carries the CRC check bits in frame 0 and the indicator bits in frames 1, 34 and 35 (used for OAM or Operations, Administration and Maintenance). The fast byte in other frames is assigned in even-/odd-frame pairs to either the EOC (Embedded Operations Channel between ATU-R and ATU-C) or to synchronization control of the bearer channels assigned to the fast buffer.
ASx and LSx: groups of bytes used for the AS (simplex bearer) channels and the LS (buplex bearer) channels.
AEX: A(S) extension bytes: bytes inserted in the transmitted ADSL frame structure to provide synchronization capacity that is shared among ASx bearer channels
LEX: L(S) Extension bytes: bytes inserted in the transmitted ADSL frame structure to provide synchronization capacity that is shared among LSx and ASx bearer channels
FEC: Forward Error Correction (calculated on the K_F bytes of the Mux data frame)
Figure illustrating
the fast byte usages

Dia 55

The figure illustrates the frame structure of the interleaved data buffer at reference point A (without FEC) and at reference point B (with FEC). Reference point C is not shown (interleaving). Note that this is more complex compared to the fast data buffer due to the FEC calculation over S frames (where S may be 1) of the superframe.
Note: The synchronization byte of the interleaved data buffer ("sync byte") carries the CRC check bits for the previous superframe in frame 0. In all other frames (1 through 67), the sync byte shall be used for synchronization control of the bearer channels assigned to the interleaved data buffer or used to carry an ADSL overhead control (AOC) channel. See also figure below.

Dia 59

Two cyclic redundancy checks (CRCs) – one for the fast data buffer and one for the interleaved data buffer – shall be generated for each superframe and transmitted in the first frame of the following superframe. Eight bits per buffer type (fast or interleaved) per superframe are allocated to the CRC check bits. These bits are computed from the k message bits (mi) using the equation:
$\mathrm{crc}(D) = M(D)\, D^8 \text{ modulo } G(D)$ where:
$M(D) = m_0 D^{k-1} + m_1 D^{k-2} + \ldots + m_{k-2} D + m_k$ is the message polynomial,
$G(D) = D^8 + D^4 + D^3 + D^2 + 1$ is the generating polynomial,
$\mathrm{crc}(D) = c_0\, D7 + c_1 D6 + \ldots + c_6\, D + c_7$ is the check polynomial
D is the delay operator.
CRC is the remainder when $M(D)\, D^8$ is divided by $G(D)$. The CRC check bits are transported in the synchronization bytes (fast and interleaved, 8 bits each) of frame 0 for each data buffer.
At the decoding side (M'(D) is received message) one will do the same division but now using $M'(D).D^8 + \mathrm{crc}(D)$. If there is no error, the remainder will be all zeros.
Note: {M'(D) D8 + crc(D)}modulo G(D) = 0 if M'(D) = M(D)

Dia 60

The figure shows a simple example of a CRC calculation. The message is 11100110 and the generator polynomial is 11001. The remainder after division is 0110 (the CRC code). First 4 zeros are appended to the message, which is equivalent to multiplying with 2^4. This is then divided (modulo 2) by the generator polynomial. The modulo 2 division is equivalent to performing the exclusive-OR operation bit-by-bit in parallel as each bit in the dividend is processed. Also, with modulo 2 arithmetic, we can perform a division into each remainder, providing the two numbers are of the same length (that is the most significant bits are both 1s. We do not consider the relative magnitudes of both numbers.

Dia 61

The circuit operates as follows. The feedback shift register (top) is cleared and the first byte in the frame is parallel-loaded into the transmit shift register (bottom). This is then shifted out to the transmission line, most significant bit first, at a rate determinated by the transmitter clock. In time synchronism with this, the same bitstream is exclusive-ORed with x3 and passed via the feedback path to the selected inputs of the feedback shift register. As each subsequent 8-bit byte is loaded into the transmit shift register and bit-serially transmitted to line, the procedure repeats. Finally, after the last byte in the frame has been output, the transmit shift register is loaded with zeros and the feedback control signal changes from 1 to 0 so that the current contents of the feedback shift register – the computed remainder – follows the frame contents onto the transmission line.

Dia 63

The binary data streams output (LSB of each byte first) from the fast and interleaved data buffers shall be scrambled separately using the following algorithm for both:

$$d'_n = d \quad \text{XOR} \quad d'_{n-18} \quad \text{XOR} \quad d'_{n-23}$$

where *dn* is the *n*-th output from the fast or interleaved buffer (i.e. input to the scrambler), and *d'n* is the *n*-th output from the corresponding scrambler.
Descrambling is done in the same way, reversing the scrambling of the data.
The scramblers are applied to the serial data streams without reference to any framing or symbol synchronization.

Dia 64

The example clearly shows a strong reduction in the transmission of long sequences of zeros or ones.

Dia 65

Block codes are memoryless codes as each output codeword depends only on the current k-bit message (e.g. parity bits may correct a single error). In contrast, with a convolutional code, the continuous stream of source bits is operated upon to produce a continuous stream of output (encoded) bits. Because of the nature of the encoding process, the sequence of source bits is said to be convolved (by applying a specific binary operation on them : e.g. shifting and adding) to produce the output bit sequence.
In ADSL, R (i.e. R_F or R_I) redundant check bytes c_0 , c_1, ..., c_{R-2}, c_{R-1} shall be appended to K (i.e. K_F or S x K_I) message bytes m_0, m_1, ..., m_{K-2}, m_{K-1} to form a Reed-Solomon codeword of size N = K + R bytes. Possible R values are : 0,2,4,6,8,10,12,14 and 16; possible S values are 1,2,4,8,16 (only for the interleaved data).

We will use a simplified example to explain the principle of a Reed-Solomon convolutional encoder/decoder (see figure). With this encoder, the three-bit shift register provides the memory and the two modulo-2 adders the convolution operation. For each bit in the input sequence, two bits are output, one from each of the two modulo-2 adders. The encoder shown is thus known as a rate 1/2(*K/N*) convolutional encoder with a constraint length of 3.
The operation is illustrated in the table (each line represents a next clock cycle). Initially. we assume that the shift register is cleared, that is, it is set to all 0s. After the first bit in the input sequence has been shifted (entered) into the shift register its contents are 001. The outputs from the two modulo-2 adders are 0 + 1 = 1 (adder 1) and 0 + 1 = 1 (adder 2). Thus, the first two output bits are 11 and these are output before the next input bit is entered into the shift register. Since the input bit was a 1, the lower branch path on the trellis diagram (see later) is followed and the output is 11, as derived. After the second input bit has been entered, the shift register contains 011. The two adder outputs are 0 + 1 = 1 (adder 1) and 1 + 1 = 0 (adder 2). Thus, the two output bits are 10 and again these are output before the next input bit is processed. Again, since the input bit was a 1, the lower branch on the trellis diagram is followed and the output is 10, as derived. Continuing, the third input bit makes the shift register contents 110 and hence the two output bits are 11; 1 + 0 = 1 (adder 1) and 1 + 0 = 1 (adder 2). Also, since the input bit was a 0, the upper branch path on the trellis diagram is followed. This process then continues.

Dia 66

The figure illustrates the tree representation of a Reed Solomon encoder. The branching points in the tree are known as nodes and the tree shows the two possible branches at each node; the upper of the two branches corresponds to a 0 input bit and the lower branch to a 1 bit. The pair of output bits corresponding to the two possible branches at each node are shown on the outside of each branch line. As we can see, with a tree diagram the number of branches in the tree doubles for each new input bit. However, the tree is repetitive after the second branch level since, after this level, there are only four unique branch nodes. These are known as *states* and are shown as *A, B, C*, and *D* in the figure. As we can see, from any one of these nodes the same pair of output bits and new node state occurs, irrespective of the position of the node in the tree. For example, from any node *C* the same pair of branch alternatives occur: 10 output and new state *A* for a 0 input, or 01 output and new state *B* for a 1 input.

Dia 67

The trellis diagram shows the outputs that result from the encoder for all possible input bit sequences. Then, for a specific input sequence, a single path through the trellis - and hence sequence of output bits – results. As an example, the lower figure shows the path through the trellis, and hence the output sequence, corresponding to the input sequence 110101....

Dia 68

The aim of the decoder is to determine the *most likely* output sequence, given a received bitstream (which may have errors) and a knowledge of the encoder used at the source. The decoding procedure is equivalent to comparing the received sequence with all the possible sequences that may be obtained with the respective encoder and then selecting the sequence that is closest to the received sequence. The Hamming distance between two codewords is the number of bits that differ between them. Therefore, when selecting the sequence that is closest to the received sequence, the Hamming distance between the received sequence and each of the possible sequences is computed, and the one with the least distance is selected. Clearly, in the limit this necessitates comparing the complete received sequence with all the possible sequences, and hence paths through the trellis. This is impractical in most cases and hence we must compromise.

Essentially, a running count is maintained of the distance between the actual received sequence and each possible sequence but, at each node in the trellis, only a single path is retained. There are a1ways two paths merging at each node and the path selected is the one with the minimum Hamming distance, the other is simply terminated. The retained paths are known as **survivor paths** and the final path selected is the one with a continuous path through the trellis with a minimum aggregate Hamming distance. This procedure is known as the **Viterbi algorithm**. The decoder, which aims to find the most likely path corresponding to the received sequence, is known as a maximum-likelihood decoder.

EXAMPLE: Assume that a message sequence of 1001110... is to be sent using the encoder shown before. From the trellis diagram for this encoder, we can deduce that this will yield a transmitted (output) sequence of:

11 01 10 11 10 00 11...

Now assume a burst error occurs so that two bits of this encoded sequence are corrupted during transmission. The received sequence is as follows:

11 01 00 11 11 00 11...

↑ ↑

Use the Viterbi algorithm to determine from this the most likely transmitted sequence.

Dia 69

Answer:

The figure (top) shows how the survivor paths are chosen. The number shown by each path merging at a node is the accumulated Hamming distance between the path followed to get to that node and the actual received sequence.

If the path chosen is that starting at the root node (branch level 0), the received sequence is 11 and the Hamming distances for the two paths are 2 for path 00 and 0 for path 11. These two distance values are added to the paths emanating from these nodes. Thus, at branch level 1, the received sequence is 01 and the two paths from node *A* have Hamming distances of 1 for path 00 and 1 for path 11. The accumulated distances are thus 2 + 1 = 3 for each path. Similarly, the two paths emanating from node *B* have Hamming distances of 0 for path 01 and 2 for path 10, and hence the accumulated distances are 0 + 0 = 0 and 0 + 2 = 2, respectively. A similar procedure is repeated at branch level 2.

At branch level 3 and onwards, however, the selection process starts. Thus, the two paths merging at node *A* (at branch level 3) have accumulated distances of 3 and 1, of which the latter is selected to be the survivor path for this node - this is shown as a bold line on the trellis diagram. A similar selection process is followed at nodes *B*, *C* and *D*. At node *C*, however, we can see that the two merging paths both have the same accumulated distance of 4. In such cases, the upper path is selected. Also, after the selection process, all subsequent distances are calculated relative to the accumulated distance associated with the selected path.

It now remains to select the most likely path and hence the output sequence. Although the decoding procedure continues, by inspection of the portion of the trellis shown, we can see that:

• only four paths have a continuous path through the trellis;
• the distance corresponding to the path *ABCABDDC* is the minimum.

Thus, this is the path that is selected, the corresponding output sequence being 11 01 10 11 10 00 11..., which corresponds to the original encoded (and hence transmitted) sequence.

Dia 70

The Reed-Solomon codewords in the interleave buffer shall be convolutionally interleaved. The interleaving depth varies, but shall always be a power of 2. Convolutional interleaving is defined by the rule: Each of the N bytes Bj0, Bj1, ..., BjN-1 in a Reed-Solomon codeword (j) are delayed by an amount that varies linearly with the byte index. More precisely, byte Bji (with index i) is delayed by (D–1) x i bytes, where D is the interleave depth. An example for N = 5, D = 2 is shown in the table, where Bji denotes the i-th byte of the j-th codeword.

Dia 71

A DMT time-domain signal has a high peak-to-average ratio (its amplitude distribution is almost Gaussian), and large values may be clipped by the digital-to-analogue converter. The error signal caused by clipping can be considered as an additive negative impulse for the time sample that was clipped. The clipping error power is almost equally distributed across all tones in the symbol in which clipping occurs. Clipping is therefore most likely to cause errors on those tones that, in anticipation of a higher received SNR, have been assigned the largest number of bits (and therefore have the densest constellations). These occasional errors can be reliably corrected by the FEC coding if the tones with the largest number of bits have been assigned to the interleave buffer. The numbers of bits and the relative gains to be used for every tone shall be calculated in the ATU-R receiver, and sent back to the ATU-C according to a defined protocol (not discussed further). The pairs of numbers are typically stored, in ascending order of frequency or tone number *i*, in a bit and gain table.

The "tone-ordered" encoding shall first assign the 8 x N_F bits from the fast data buffer to the tones with the smallest number of bits assigned to them, and then the 8 x N_I bits from the interleave data buffer to the remaining tones. All tones shall be encoded with the number of bits assigned to them; one tone may therefore have a mixture of bits from the fast and interleaved buffers. The ordered bit table b'_i shall be based on the original bit table b_i as follows:

For k = 0 to 15 { From the bit table, find the set of all i with the number of bits per tone b_i = k, Assign b_i to the ordered bit allocation table in ascending order of I }

A complementary de-ordering procedure should be performed in the ATU-R receiver. It is not necessary, however, to send the results of the ordering process to the receiver because the bit table was originally generated in the ATU-R, and therefore that table has all the information necessary to perform the de-ordering.

Dia 73

The figure shows the overall architecture of an ATM (Asynchronous Transfer Mode) based DSL access network. ATM is used between the ATU-R and ATU-C and further it is also

used to interconnect the DSLAMs to BRAS (Broadband Remote Access Servers). This is sometimes called the aggregation network. ATM flows will start at the BRAS and terminate in the ATU-R (or the other way around). Note that ATM will typically transport PPP frames that are filled with IP packets.
An ADSL system transporting ATM shall support the single latency mode* at all integer multiples of 32 kbit/s up to 6.144 Mbit/s downstream and up to 640 kbit/s upstream. For single latency, ATM data shall be mapped to bearer channel AS0 in the downstream direction and to bearer channel LS0 in the upstream direction. Single latency is defined as all payload data passing through a single latency path .
ADSL systems transporting ATM shall support bearer channel AS0 downstream and bearer channel LS0 upstream, with each of these bearer channels independently allocable to a particular latency path as selected by the ATU-C at start-up. Therefore, support of dual latency is optional for both downstream and upstream. If downstream ATM data are transmitted through a single latency path (i.e. "fast" only or "interleaved" only), only bearer channel AS0 shall be used, and it shall be allocated to the appropriate latency path. If downstream ATM data are transmitted through both latency paths (i.e. "fast" and "interleaved"), only bearer channels AS0 and AS1 shall be used, and they shall be allocated to different latency paths. Similarly, if upstream ATM data are transmitted through a single latency path (i.e. "fast" only or "interleaved" only), only bearer channel LS0 shall be used and it shall be allocated to the appropriate latency path. The choice of the fast or interleaved path may be made independently of the choice for the downstream data. If upstream ATM data are transmitted through both latency paths (i.e. "fast“ and "interleaved"), only bearer channels LS0 and LS1 shall be used and they shall be allocated to different latency paths. Bearer channel AS0 shall support the transport of data at all integer multiples of 32 kbit/s from 32 kbit/s to 6.144 Mbit/s. Bearer channel LS0 shall support all integer multiples of 32 kbit/s from 32 kbit/s to 640 kbit/s. Support for data rates based on non-integer multiples of 32 kbit/s is also optional.
* For ATM systems, the channelization of different payloads is embedded within the ATM data stream using different Virtual Paths and/or Virtual Channels. Therefore the basic requirements for ATM are for only one ADSL bearer channel downstream and only one ADSL bearer channel upstream.
ATM is using the same principle as MPLS : labels (VPI or Virtual Path Identifiers and VCI or Virtual Channel Identifiers that are part of the header) are used to define the path in the network. A path will be specified by using translation tables (translating the VPI and VCI) in each node on the path. An ATM packet (called a cell) has a fixed length of 5 bytes header overhead and 48 bytes payload.

Dia 75

This figure gives an overview of some recent evolutions in xDSL technology. Some of the goals are: increase the available bit rate, make the streams more symmetric and try to keep the distance bridged as large as possible.

Dia 76

VDSL transceivers shall use Frequency Division Duplexing (FDD) in separating upstream and downstream transmission. VDSL systems use a four-band plan that starts at 138 kHz and extends up to 12 MHz. The four frequency bands are denoted as DS1, US1, DS2, and US2(DS: downstream; US upstream). Frequency spacing of the DMT subbands (or tones)

is 4.3125 kHZ and in total there are 2n+8 tones possible (n=0,1,2,3,4) or 256, 512, 1024, 2048 or 4096 tones.

Dia 78

CATV networks (Community Antenna TeleVision) have for a long time been solely deployed for the distribution of television programs to the subscribers, coaxial cable was used as medium of transmission. Present CATV networks are therefore optimized for unidirectional transmission of distribution services.
The network basically consists of a tree and branch coaxial cable in order to distribute TV-signals (typical bandwidth / channel: 6 to 8 MHz) from the “Head End”to the homes.
In Europe the CATV program channels are typically located in the VHF band (from 47 - 68 MHz) next to the FM radio channels (from 87,5 - 104 MHz) and in the UHF band (from 174 - 230 MHz and from 470 - 960 MHz).
The network is unidirectional and distributive. There are typically a large number of users connected (in the past 50000 or more) to one Head End (HE). In the Head End, the different TV-signals are captured (from satellite, optical fiber, terrestrial wireless connections) or locally generated. They are combined on the correct carrier frequencies using FDM (Frequency Division Multiplexing).
The household penetration is largely depending on the country. Some examples are : Belgium 90%, The Netherlands 87%, Germany 53%, Denmark 55%, Sweden 51%, Portugal 23%, UK 13%, France 12%, Spain 4%, Italy <1%, Europe (EU15) 30% [CATV subscribers / households in %, figures from 2000].

Dia 79

A typical tree and branch topology of the CATV network is shown in the figure. We observe that amplifiers are used in order to overcome the transmission and splitting losses. This results however in a unidirectional system because the amplifiers are unidirectional.
There are 3 parts in the network : Trunk Cable (with trunk amplifiers), Feeder Cable (with feeder amplifiers) and Drop Cable.
Note that the coaxial network is a shared medium access network, this means that the available bandwidth needs to be shared by all customers connected to one “branch” of the network. Because for TV-distribution all customers receive the same signal, this is not a limitation.

Dia 80

When adapting the coax network for the provision of interactive multimedia services (such as telephony, data communication, internet access, VOD) it is predominant to:
go from analog to digital (but use overlay structure to support current analog TV-services)
provide bidirectional transmission (a split of the spectrum in upstream and downstream is required to allow bidirectional amplification)
decrease the distance covered by coax cable in order to reduce the sharing of the same transmission medium
improve the bandwidth usage by implementing advanced modulation techniques
connect the network to a backbone network (telephony, internet).
As already discussed in the TP network, advanced modulation techniques are incorporated in modems to allow a more efficient use of the available bandwidth. Next to that, the frequency band used on the coax network, will also be expanded by the introduction of

improved amplifiers (upto 1 GHz). Upgrading the capacity available for each user, will be done by gradually deploying optical fibre in the network and decreasing the amount of users connected to the same network branch. Splitting up the available spectrum, as shown on the next slides, enables bi-directionality in the network.
Next to these measures to allow broadband services, interactive services will only be possible by introducing switching/routing capability in the network. This can be obtained in the head end by interconnection to a switched (broadband) network
It is clear that although the operators are eager to upgrade their coax network, the current analog TV-channels should still be available on the network (not everybody will switch at the same time to interactive services !).

Dia 81

This figures illustrates a coax cable network that is adapted to support interactive multimedia services. One observes the bidirectional amplifiers (for bidirectional transmission), the optical fiber and the interconnection with PSTN and Internet.
This network is sometimes called a HFC or Hybrid Fiber Coax network.
The number of users typically sharing the coaxial tree and branch is in the range of 500 to 2000.

Dia 82

The spectrum used on the coaxial cable has to be split for upstream (from user to network) and downstream (from network to user) traffic. The exact split of the available coax spectrum and the allocation of the frequency plan depend on operator and country (regulation).
Generally speaking we can state that the upstream traffic will be allocated in the lower part of the frequency band (between 5 MHz and 25 - 65 MHz). The upstream direction will carry the voice telephony signal but also packets sent to the Internet or commands to request video services.
The major part of the spectrum is allocated for downstream traffic. The analogue TV-channels are allocated in 6-8 MHz channels. Exploitation of the coaxial frequency spectrum is foreseen in the near future up to the GHz region. Upper regions in the frequency spectrum are typically allocated for new interactive services. Please note that once certain analogue TV-channels are no longer used in the future, these frequency bands can be used for digital TV or other types of services.
One of the key advantages of the coax access network is the ability to carry multiple types of information in multiple formats shared by a scalable number of users. This means that when using appropriate modulation and compression techniques, a large range of services can be offered to the users.

Dia 83

This figure illustrates in more detail the usage of the spectrum in the current CATV network. In order to introduce iDTV, a number of analog channels were converted in to digital channels (not shown on the figure).

Dia 84

A recent standard that supports digital bidirectional access on coax plants is Docsis (Data Over Cable System Interface Specification). This is a US-standard that has been adapted to

the European cable networks (Euro Docsis). The standard is specifying the physical layer (different for US and Europe) and the data link layer (a.o. Medium Access Control protocol on the shared upstream coax part).

Dia 85

The general architecture is shown in the top figure : A CMTS (Cable Modem Termination System) will be connected to the public network (left interface) and to the HFC network on the right side. The network consists of a fiber part and a tree and branch coax part. At the customer premises, the analog TV-signals are removed from the signal and the digital signals are decoded in the Cable Modem (CM). The interface that is provided is Ethernet (10 or 100 Mbit/s) or Universal Serial Bus or USB (12 Mbit/s or USB with 480 Mbit/s) or Peripheral Component Interface or PCI (132 MByte/s or 1 GByte/s). WiFi (IEEE 802.11) is also an option.
Some general characteristics of the Euro Docsis system are listed in the table. The downstream direction will make use of a frequency range up to about 860 MHz (starting from about 100 MHz). This frequency range will be split in bands of 8 MHz (=analog TV-channel separation) which will be modulated using 64 or 256 QAM (Quadrature Amplitude Modulation). The resulting bitrates in a single 8 MHz channel range approximately between 40 and 50 Mbit/s. MPEG is used in the downstream direction. Note that FM radio is also distributed over the coax (therefore the downstream frequency starts at 108 MHz).
The upstream direction is more complicated because a number of users may send some information at the same time to the head end. This will require a medium access control protocol (TDMA or Time Division Multiple Access based). The available frequency range is from 5 to 65 MHz* and it is split in frequency bands in the range of 0.2 to 3.2 MHz. Depending on the modulation format (16-QAM or QPSK Quadrature Phase Shift Keying) and the bandwidth of the frequency bands, bitrates ranging from 0.32 to 10 Mbit/s are available.
* The frequency spectrum between 5 and 20 MHz will suffer from disturbances and will not be used in general.

Dia 86

The various protocols that are used in the DOCSIS standard all have specific functions. The relationships of these protocols are shown in the diagram labeled "DOCSIS Protocols". These protocols are the basis of operation for any cable modem. The lower four layers are specific to the cable data network (HFC plant) and are the foundation of communication between any cable modem (CM) and the cable modem termination system (CMTS). The upper layers are protocols that are carried over the communication path established by the lower layers. Well known IP services such as email, web, file transfer, and network news are all presented as TCP/IP traffic. Other protocols such as IPX/SPX, AppleTalk, and NetBeui are possible however since a vast majority of the internet is based on TCP/IP, they are not shown here.
The lower four layers of DOCSIS are:
PHY - physical layer. This layer defines the modulation schemes used on a HFC plant. This layer is responsible for translating the raw signal found on the actual coaxial cable to and from digital information used by the upper layers in the protocol stack. The modulation schemes present on the coaxial cable can be 64 or 256 quadrature amplitude modulation

(64- or 256-QAM) for the downstream and quadrature phase shift keying or 16-QAM for the upstream.
MPEG-2 - transmission convergence layer. All data that is present on the downstream is encapsulated into MPEG-2 frames. These frames can contain the actual video and audio data that is typically decoded and presented as TV image and sound. They can also contain data that is decoded and presented as information available for computer usage (i.e. the internet). Since all data... voice, audio, internet... is encoded into MPEG-2 frames, it is quite feasible for a cable operator to multiplex all signals onto one cable. This is inherently how cable internet access works.
MAC - media access control layer. All data that is present on the upstream is managed by this section of the DOCSIS protocol stack. Since a HFC plant is similar to an Ethernet network in that all communication devices are "connected" to the same cable, it is imperative that an orderly process exist whereby the CMTS can tell the CM when to transmit data and for how long. The MAC layer is the means of coordinating upstream traffic from the CM to the CMTS.
BPI - data link encryption. Since the cable network is a shared medium, there must be a method to protect user data from malicious usage. The DOCSIS standard defines Baseline Privacy Interface as this method.

Dia 87

The upstream channel is characterized by many transmitters (CMs) and one receiver (CMTS). Time in the upstream channel is slotted, providing for Time Division Multiple Access at regulated time ticks. The CMTS provides the time reference and controls the allowed usage for each interval. Intervals may be granted for transmissions by particular CMs, or for contention by all CMs. CMs may contend to request transmission time. Some highlights of the MAC protocol v1.1 include:
- Bandwidth allocation controlled by CMTS
- A stream of mini-slots in the upstream
- Dynamic mix of contention- and reservation-based upstream transmit opportunities
- Bandwidth efficiency through support of variable-length packets
- Extensions provided for future support of ATM or other Data PDU
- Quality-of-service features
- Support for a wide range of data rates.

Dia 88

Upstream bandwidth allocation

The upstream channel is modeled as a stream of mini-slots. The CMTS must generate the time reference for identifying these slots. It must also control access to these slots by the cable modems. For example, it may grant some number of contiguous slots to a CM for it to transmit a data PDU. The CM must time its transmission so that the CMTS receives it in the time reference specified. The basic mechanism for assigning bandwidth management is the allocation MAP. The allocation MAP is a MAC management message transmitted by the CMTS on the downstream channel which describes, for some interval, the uses to which the upstream mini-slots must be put. A given MAP may describe some slots as grants for particular stations to transmit data in, other slots as available for contention transmission, and other slots as an opportunity for new stations to join the link. Many different scheduling algorithms may be implemented for the CMTS by different vendors;

this specification does not mandate a particular algorithm. Instead, it describes the protocol elements by which a bandwidth is requested and granted.

The allocation MAP MAC management message

The allocation MAP is a varying-length MAC Management message that is transmitted by the CMTS to define transmission opportunities on the upstream channel. It includes a fixed-length header followed by a variable number of information elements (IEs). Each information element defines the allowed usage for a range of mini-slots.

Dia 89

Quality of Services

The DOCSIS RFI v1.1 provide the following five classes of service for the traffic.
- Unsolicited Grant Service (UGS)
- Real-Time Polling Service (rtPS)
- Unsolicited Grant Service with Activity Detection (USGAD)
- Non-Real time polling service (nrtPS)
- Best effort Service (BE)

UGS support isochronous traffic such as CBR type traffic. The CMTS gives periodic grant for this traffic. The rtPS provide the rt-VBR type service in ATM but uses the polling mechanism to support the real time traffic. USG-AD is a new service in the DOCSIS version 1.1. It uses the CM activity detection technique and gives the grant when the CM is active. This service can support voice traffic with silence suppression. nrtPS and BE are similar to non-real time VBR and UBR services in ATM. To provide these QoS to the incoming traffic, following concepts are added in the DOCSIS version 1.1.
- Packet Classification and Flow Identification
- Service Flow QoS Scheduling
- Dynamic Service Establishment
- Fragmentation

Dia 90

The figure illustrates that a MAP PDU is received by the CM and in this way the US channel is defined. The US channel MAP typically starts with a number of minislots reserved for access request (in contention mode). A number of blocks follow where different CMs are allowed to send information upstreams.

Dia 91

The principle topology of the network is shown in the figure.
Originally the CATV network consisted of head-ends (where the video signals were captured and put on the coax) of the different cable operators and a tree and branch coaxial network. There were a large number of head-ends because of the large number of CATV operators in Flanders.
Each coax tree and branch network is currently split in islands (connected to a node, dots on the figure) and these nodes are connected via a fiber ring to the original headend (triangle on the figure). The nodes are serving about 1000 CATV subscribers (between 125 and 2000) via the existing upgraded coax plant. The headends themselves are interconnected to the new switches (5 in total, represented by squares on the figure) by the use of 8 secondary rings connected to a primary ring.
The actual topology is shown on the next figure.

Dia 92

Topology of the Telenet Operations network : primary ring and secondary rings.

Dia 93

Fiber network of Telenet.

Dia 96

The figure illustrates the usage of fiber in access networks. The FTTC and HFC scenario's were mentioned already during the description of xDSL and coax networks.
The two lower figures illustrate FTTH (Fiber To The Home) architectures. Two options are: the PON or Passive Optical Network solution (operating in a point-to-multipoint mode) and a point-to-point based solution. They will be discussed in the next slides

Dia 97

In an optical access network, the Central Office (OC) contains an optical line terminal (OLT) which provides the network-interface, and this OLT is connected to one or more optical network units or terminals (ONU or ONT) at the user-side. Dependent on the used architecture (either active or passive) and the transmission protocol (either Ethernet or ATM), we can distinguish numerous topologies.
Passive Optical Networks or PONs are point-to-multipoint connections, made up of fibre optic cabling, of passive splitters and couplers that distribute an optical signal through a branched "tree" topology to connectors that terminate each fibre segment. For PON architectures, there is the choice between either Ethernet or ATM. Two groups are working on these standards: IEEE 802.3ah (EFM: Ethernet in the First Mile, also promoted by Ethernet in the First Mile Alliance or EFMA) and FSAN (Full Service Access Network, backed by the ITU).
EPON or Ethernet PON (IEEE 802.3ah).
Recently, the possibility of a 10-Gb/s EPON standard (known as next generation EPON or NGEPON) is being proposed to the IEEE.
BPON or Broadband PON (ITU–T G.983).
Originally: APON or ATM PON. The initial PON specifications defined by the FSAN committee used ATM as their layer 2 signalling protocol. Use of the term APON led users to believe that only ATM services could be provided to end-users, so the FSAN decided to broaden the name to Broadband PON (BPON). BPON systems offer numerous broadband services including Ethernet access and video distribution.
GPON or Gigabit PON (ITU–T G.984).
In 2001 the FSAN group initiated a new effort for standardizing PON networks operating at bit rates of above 1 Gb/s. Apart from the need to support higher bit rates, the overall protocol has been opened for re-consideration and the sought solution should be the most optimal and efficient in terms of support for multiple services, OAM&P (Operations, Administration, Maintenance and Provisioning) functionality and scalability.
Instead of having a PON, it is also possible to deploy an active network (with point-to-point connections), which looks very similar to a PON, however with some important differences. The most fundamental one is to replace the passive, unmanageable splitters in the field by an active node. An important consequence is that a power line between the CO and the active node will be necessary. For point-to-point fiber topologies, the standard

technology is (Active) Ethernet, also defined by the IEEE (IEEE 802.3ah) and ITU (ITU-T G.985).
More details about the different standards can be found on the next slide.

Dia 98

The table gives an overview of the most important technologies and related standards for optical access networks.

Dia 99

The graphs illustrate the recent deployments of FTTx.
Roughly, the geographical distribution of the different protocols (see above) is as follows:
Europa: Point-to-Point Ethernet (IEEE 802.3ah)
Japan: Point-to-Multipoint Ethernet or EPON (IEEE 802.3ah)
USA: BPON/GPON (ITU-T G.983/G.984)

Dia 100

The top figure shows a typical PON deployment scenario where a fiber is leaving the OLT (central office side) and is split in 4 fibers. The first part uses buried optical fiber cable, the second part (closer to the user) uses overhead cable (connected to poles).
The bottom figure illustrates a possible architecture including the supported services and interfaces. Note that the PON is using CWDM (Coarse Wavelength Division Multiplexing) using the 1500 nm band for downstream traffic and the 1300 nm band for upstream traffic. An overlay network at 1550 nm is distributing CATV video signals.

Dia 102

Ranging: measuring the distance between the cable modems (at home) and the CMTS. This is important to schedule the upstream traffic (one has to take into account the delays in order to avoid collisions of upstream packets).

Dia 104

Data is broadcast downstream from the OLT to multiple ONUs in variable-length packets of up to 1,518 bytes, according to the IEEE 802.3 protocol. Each packet carries a header that uniquely identifies it as data intended for ONU–1, ONU–2, or ONU–3. In addition, some packets may be intended for all of the ONUs (broadcast packets) or a particular group of ONUs (multicast packets). At the splitter, the traffic is divided into three separate signals, each carrying all of the ONU–specific packets. When the data reaches the ONU, it accepts the packets that are intended for it and discards the packets that are intended for other ONUs.

EPON Frame Formats

The downstream traffic is segmented into fixed-interval frames, each of which carries multiple variable-length packets. Clocking information, in the form of a synchronization marker, is included at the beginning of each frame. The synchronization marker is a one-byte code that is transmitted every 2 ms to synchronize the ONUs with the OLT.
Each variable-length packet is addressed to a specific ONU as indicated by the numbers, 1 through N. The packets are formatted according to the IEEE 802.3 standard and are transmitted downstream at 1 Gbps. The expanded view of one variable-length packet shows the header, the variable-length payload, and the error-detection field.

Dia 105

Upstream traffic is managed by utilizing TDM technology, in which transmission time slots are dedicated to the ONUs. The time slots are synchronized so that upstream packets from the ONUs do not interfere with each other once the data is coupled onto the common fiber. For example, ONU–1 transmits packet 1 in the first time slot, ONU–2 transmits packet 2 in a second non-overlapping time slot, and ONU–3 transmits packet 3 in a third non-overlapping time slot.

EPON Frame Formats

The upstream traffic is segmented into frames, and each frame is further segmented into ONU–specific time slots. The upstream frames are formed by a continuous transmission interval of 2 ms. A frame header identifies the start of each upstream frame.
The ONU–specific time slots are transmission intervals within each upstream frame that are dedicated to the transmission of variable-length packets from specific ONUs. Each ONU has a dedicated time slot within each upstream frame. For example, each upstream frame is divided into N time slots, with each time slot corresponding to its respective ONU, 1 through N.
The TDM controller for each ONU, in conjunction with timing information from the OLT, controls the upstream transmission timing of the variable-length packets within the dedicated time slots. The figure shows an expanded view of the ONU–specific time slot (dedicated to ONU–4) that includes two variable-length packets and some time-slot overhead. The time-slot overhead includes a guard band, timing indicators, and signal power indicators. When there is no traffic to transmit from the ONU, a time slot may be filled with an idle signal.

Dia 106

The IEEE 802.3ah Task Force is developing the so-called multipoint control protocol (MPCP), which arbitrates channel access among central office and subscribers. MPCP is used to dynamically assign the upstream bandwidth (subscriber to service provider), which is the key challenge in access protocol design for EPONs. Note that MPCP does not specify any particular dynamic bandwidth allocation (DBA) algorithm. Instead, it is intended to facilitate the implementation of DBA algorithms.
The MPCP arbitration mechanism is used to dynamically assign nonoverlapping upstream transmission windows (time slots) to each ONU. Besides auto-discovery, registration, and ranging (RTT computation) operations for newly added ONUs, MPCP provides the signaling infrastructure (control plane) for coordinating data transmissions from the ONUs to the OLT. The basic idea is that the upstream bandwidth is divided into bandwidth units via TDM. These units are assigned to the ONUs as determined by the OLT according to the DBA algorithm in use. The OLT has control over the assignment of these units of bandwidth. These units can be assigned on the fly as needed or can be reserved in advance. For efficiency reasons, any reserved units or fraction of units of bandwidth that go unused can in general be re-assigned on the fly by the OLT to other ONUs that could make use of it.

Dia 107

MPCP uses two types of messages to facilitate arbitration: REPORT and GATE. Each ONU has a set of queues, possibly prioritized, holding Ethernet frames ready for upstream transmission to the OLT. The REPORT message is used by an ONU to report bandwidth

requirements (typically in the form of queue occupancies) to the OLT. A REPORT message can support the reporting of up to 13 queue occupancies of the corresponding ONU. Upon receiving a REPORT message, the OLT passes it to the DBA algorithm module. The DBA module calculates the upstream transmission schedule of all ONUs such that channel collisions are avoided. After executing the DBA algorithm, the OLT transmits GATE messages to issue transmission grants. Each GATE message can support up to four transmission grants. Each transmission grant contains the transmission start time and transmission length of the corresponding ONU. Each ONU updates its local clock using the timestamp contained in each received transmission grant. Thus, each ONU is able to acquire and maintain global synchronization. The transmission start time is expressed as an absolute timestamp according to this global synchronization. Each ONU sends backlogged Ethernet frames during its granted transmission window using its local intra-ONU scheduler. The intra-ONU scheduler schedules the packet transmission from the various local queues. The transmission window may comprise multiple Ethernet frames; packet fragmentation is not allowed. As a consequence, if the next frame does not fit into the current transmission window, it has to be deferred to the next granted transmission window.

Dia 108

Overview of MPCP control messages.

Dia 109

The considered DBA algorithms can be used in the DBA module of the above described MPCP arbitration mechanism to calculate the collision-free upstream transmission schedule of ONUs and generate GATE messages accordingly.
We categorize the DBA algorithms for EPONs into algorithms with statistical multiplexing and algorithms with quality of service (QoS) assurances. The latter are further subdivided into algorithms with absolute and relative QoS assurances.
IPACT (and also the control theoretic extension) is elaborated in more detail in the next slides.
More detailed information about the other DBA algorithms can be found in:
M. P. McGarry, M. Maier, M. Reisslein, “Ethernet PONs: a Survey of Dynamic Bandwidth Allocation (DBA) Algorithms”, IEEE Communications Magazine, Aug. 2004.

Dia 110

We give a high-level overview of the proposed algorithm. For simplicity of illustration, we will consider a system with only three ONUs.
1. Let us imagine that at some moment of time *t0* the OLT knows exactly how many bytes are waiting in each ONU’s buffer and the Round-Trip Time (RTT) to each ONU. The OLT keeps this data in a polling table shown in Figure a. At time *t0* , the OLT sends a control message to ONU1, allowing it to send 6000 bytes (see Figure a). We will call such a message a *Grant*. Since, in the downstream direction, the OLT broadcasts data to all ONUs, the Grant should contain the ID of the destination ONU, as well as the size of the granted window (in bytes).
2. Upon receiving the Grant from the OLT, ONU1 starts sending its data up to the size of the granted window (Figure b). In our example – up to 6000 bytes. At the same time, the ONU keeps receiving new data packets from its user. At the end of its transmission

window, ONU1 will generate its own control message (*Request*). The Request sent by ONU1 tells the OLT how many bytes were in ONU1's buffer at the moment when the Request was generated. In our case there were 550 bytes.
3. Even before the OLT received a reply from ONU1, it knows when the last bit of ONU1's transmission will arrive. This is how the OLT calculates this:
(a) the first bit will arrive exactly after the RTT time. The RTT in our calculation includes the actual round-trip time, Grant processing time, Request generating time, and a preamble for the OLT to perform bit- and byte-alignment on received data, i.e., it is exactly the time interval between sending a Grant to an ONU and receiving data from the same ONU. (b) since the OLT knows how many bytes (bits) it has authorized ONU1 to send, it knows when the last bit from ONU1 will arrive. Then, knowing RTT for ONU2, the OLT can schedule a Grant to ONU2 such that first bit from ONU2 will arrive with a small guard interval after the last bit from ONU1 (Figure 5.b). The guard intervals provide protection for fluctuations of round-trip time and control message processing time of various ONUs. Additionally, the OLT receiver needs some time to readjust its sensitivity due to the fact that signals from different ONUs may have different power levels because ONUs are located at different distances from the OLT (far-near problem).
4. After some time, the data from ONU1 arrives. At the end of the transmission from ONU1, there is a new Request that contains information of how many bytes were in ONU1's buffer just prior to the Request transmission. The OLT will use this information to update its polling table (see Figure c). By keeping track of times when Grants are sent out and data is received, the OLT constantly updates the RTT entries for the corresponding ONUs.
5. Similarly to Step 4, the OLT can calculate the time when the last bit from ONU2 will arrive. Hence, it will know when to send the Grant to ONU3 so that its data is tailed to the end of ONU2's data. After some more time, the data from ONU2 will arrive. The OLT will again update its table, this time the entry for the ONU2 (see Figure d).
If an ONU emptied its buffer completely, it will report 0 bytes back to the OLT. Correspondingly, in the next cycle the ONU will be granted 0 bytes, i.e., it will be allowed to
send a new request, but no data. Note that the OLT's receive channel is almost 100 % utilized (Requests and guard times consume a small amount of bandwidth). Idle ONUs (without data to send) are not given transmission windows. That leads to a shortened cycle time, which in turns results in more frequent polling of active ONUs. As it is clear from the description above, there is no need to synchronize the ONUs to a common reference clock (as traditionally done in TDMA schemes). Every ONU executes the same procedure driven by the Grant messages received from the OLT. The entire scheduling and bandwidth allocation algorithm is located in the OLT. Thus, it is easy to adaptively change the scheduling at run-time based on some network conditions; the ONUs do not need to negotiate or acknowledge new parameters, nor do they need to switch to new settings synchronously.
If the OLT authorizes each ONU to send its entire buffer contents in one transmission, ONUs with high data volume could monopolize the entire bandwidth. To avoid this, the OLT will limit the maximum transmission size. Thus, every ONU will get a Grant to send as many bytes as it has requested in a previous cycle, but no more than some maximum limit (maximum transmission window size). There could be various schemes for specifying

the limit. It can be fixed, say, based on a Service Level Agreement (SLA) for each ONU, or dynamic - based on average network load.

Dia 111

Control Theoretic Extension of IPACT.