ISDN: Theory & Practice

routed by Mr.Goblins [2002]
updated by Mr. Vesa Kosonen [2007]

Basic Concepts

In this section you will find out:

- The definition of ISDN
- Basic Rate and Primary Rate ISDN
- What an ISDN line provides to the user
- How ISDN provides services to the user
- What the B and D channels are used for

What is ISDN?

Integrated **S**ervices **D**igital **N**etwork is a set of digital transmission protocols defined by the international standards body for telecommunications, the ITU-T (previously called the CCITT). These protocols are accepted as standards by virtually every telecommunications carrier all over the world.

ISDN complements the traditional telephone system so that a single pair of telephone wires is capable of carrying voice and data simultaneously. It is a fully digital network where all devices and applications present themselves in a digital form.

The essential difference between ISDN and the conventional telephone system is that it is digital not analogue. Information travels as bits rather than as waves. In addition, it also allows multiple streams of these bits to occupy the same connection, providing the user with greater versatility of services.

What are the applications for ISDN?

The Integrated Services Digital Network uses the twisted-pair copper telephone line that would traditionally carry only one voice connection. ISDN can carry more than one connection over this wire at the same time, and at greater speed. Applications include telecommuting; simultaneous voice, fax, data and e-mail; inexpensive videoconferencing; remote broadcasting and high quality audio transmission.

ISDN handles all types of information - voice, data, studio-quality sound, static and moving images. They are all digitised, and transmitted at high speed.

ISDN can handle many devices and many telephone numbers on the same line. Up to eight separate telephones, fax machines or computers can be linked to a single Basic Rate ISDN connection and have different phone numbers assigned to them. (We'll explain Basic Rate in the following pages).

A Basic Rate ISDN line can support up to two calls at the same time. Any combination of voice, fax or PC connections can take place at the same time, through the same ISDN line.

From a digital ISDN telephone you can place a call to an analogue telephone on the PSTN (Public Switched Telephone Network) and vice-versa. Both networks are interconnected by the network carrier in a way similar to the connection between the mobile phone network and the analogue phone network. For the user, it is completely transparent whether he is calling a GSM telephone, a conventional telephone or an ISDN digital telephone.

What does ISDN give you? - 1

There are two forms in which ISDN is supplied.

Basic Rate

- Access to the network is called **Basic Rate Access** (BRA).
- It is provided through a **Basic Rate Interface** (BRI).
- This kind of interface is also called an $\mathbf{S_0}$ Interface.
- There are **two channels** that you can use.

Primary Rate

- Access to the network is called **Primary Rate Access** (PRA).
- It is provided through a **Primary Rate Interface** (PRI).
- This kind of interface is also called an $\mathbf{S_2}$ Interface.
- There are either **30 channels** (most of the world) or 23 channels (North America, Japan) that you can use.

What does ISDN give you? – 2

Basic Rate

- The total data rate across this interface **144 000** bits per second.
- This bit rate was chosen because the wiring already installed by the telephone companies under the streets can carry baseband (digital) transmission at this speed.

Primary Rate

- The total data rate across this type of interface is **2 048 000** bits per second **(2Mbits per second)** in Europe. In North America and some other countries the total data rate is **1 536 000** bits per second **(1.5Mbits per second)**
- This kind of access requires the installation of a high-speed line to the customer premises.

Normally, Basic Rate would be for domestic use, telecommuters or smaller remote offices. Primary Rate would typically be used for large remote access servers, fax servers or PBXs in medium sized or large offices. For instance, most ISPs (Internet Services Providers) use PRI lines to provide dial-in analogue and ISDN connections for their subscribers.

ISDN Services

There are two kinds of services provided by ISDN.

Network services

- Network services carry the interactions between the user and the network
- For example: setting up calls and disconnecting them

Bearer services

- Bearer services carry data between two users
- For example: voice or fax information encoded as a bit stream

Network Services

- Network Services define how the user and the network interact with each other in order to manage calls.
- The user can use Network Services to request the network to perform functions such as making and clearing calls, transferring calls to another user, and so on.
- This activity is known as signalling.

Bearer Services

- Bearer services carry the call activity that the user is performing at any given moment.
- This includes voice calls, fax and modem calls, and connections to the Internet.
- Broadly speaking, there are two forms of bearer service.
    - **Structured Data -** the information passing over the bearer service is in a format that is understood by the network.
      Voice is an example of structured data. Because the network knows that the connection carrying voice, it can convert the data into an analogue signal in the event that the call is connected to an ordinary analogue phone.
    - **Unstructured Data** - the format of the information is not understood by the network, but is understood by the two users at either end of the service.

Summary

- Bearer services provide connections between users across the network.
- Network services provide control and signalling between the user and the network.
- ISDN is provided either as Basic rate or as Primary rate.
- Basic rate provides 2 user channels.
- Primary rate provides either 30 or 23 user channels.

Q/A

1. What is the essential difference between ISDN and the ordinary telephone system (PSTN)?

The essential difference between ISDN and the ordinary PSTN system is that ISDN is digital, whereas the PSTN is analogue.

ISDN signals pass as a stream of bits. It is from this digital basis that most of the benefits of ISDN arise: speed, multi-channeling, and the ability to carry many types of calls.

2. In addition to voice, name at least two other services which can be carried over ISDN.

In addition to voice, ISDN can also handle fax, computer data and e-mail, videoconferencing, and audio transmission.

3. Name the two forms in which an ISDN service can be provided to the user.

ISDN can be supplied to the user in two forms:

Basic Rate - Two 'B' channels

Typical usage: domestic use, home-workers, remote offices

and

Primary Rate - 30 'B' channels (23 in North America)

Typical usage: corporate use, dial-in remote access servers, Internet service providers, digital PABXs

B & D channels

In this section you will learn

- What the B and D channels do
- How the B and D channels share the line
- How the structures of BRI and PRI compare
- What Bearer Capabilities are and on which channel they reside
- How to vary the the amount of bandwidth you have available to you

What does a B channel do?

The B channel carries ISDN Bearer Services across the network and so carries the content of call (the voice, fax or data) between users.

The B channel is a neutral conduit for bits and carries data at 64 000 bits per second (56 000 bits per second in some North American networks).

The ISDN does not need to know what the bits represent. The job of the network is to accept a stream of bits supplied by one user at one end of the B channel and to deliver them to the other user at the opposite end of the channel.

Within an interface, the B channels are numbered. In a Basic Rate Interface they are numbered 1 & 2; in a Primary Rate Interface, they are numbered 1 to 30 (or 23 in North America). When two users are connected, there is no relationship between the channel numbers used at each end. You might have one user's B channel number 17 connected with the other user's B channel number 2. The ISDN is responsible for managing this relationship.

Notice that channel number 17 would only be possible on a PRI, while channel number 2 is possible on both a BRI and a PRI. ISDN does not restrict the interconnection of B channels between the two kinds of interface.

What does the D channel do?

The D channel carries the ISDN Network Services between the user and the network. It maintains the user's relationship with the network.

This includes:

- the requests and responses used when you make or receive a call
- call progress messages
- messages informing you that the called party has closed the call
- error messages telling you why a call has not been established for you

The D channel operates at 16 000 bits per second in a BRI and at 64 000 bits per second in a PRI.

B & D channel characteristics

An ISDN channel has two and only two ends. B channels terminate at a user. A B channel can therefore connect two and only two users. A B channel cannot be Y-shaped. B channels are therefore described as **end-to-end**.

In the case of the D channel one end is with the user. The other end is in the network.

A D channel is not end-to-end. You cannot normally, therefore, use a D channel to carry data between two users.

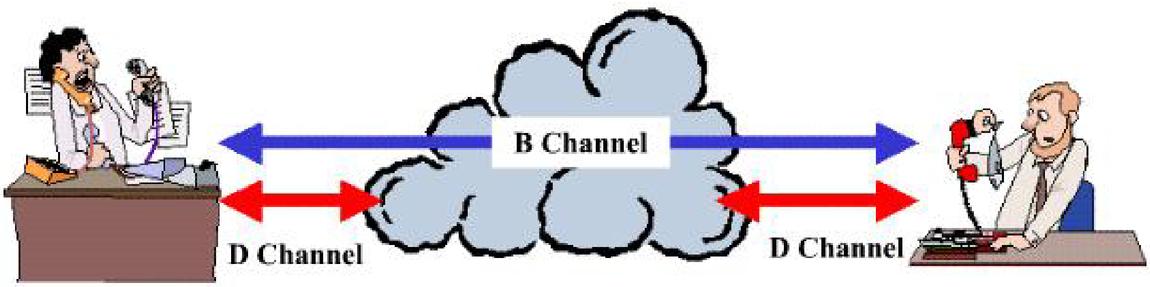

Notice how the D channels (the red lines) do not pass through the network. Notice also how each user has only one D channel and it is not connected in any way with the D channel of the other user.

The B channel (the blue line) passes directly across the network.

How B & D channels share the line: Basic rate

The two B channels and the one D channel that make up a Basic Rate ISDN line are assembled together within the interface using a technique called **Time Division Multiplexing**.

It works like this. Imagine a clock with the second hand spinning. For a portion of the arc described by the second hand, the interface is carrying data for the D channel; for another portion of the arc, the interface is carrying data for a B channel.

How B & D channels share the line: Primary rate

A EuroISDN Primary Rate Interface contains **30 B Channels and one D channel**. In North America and some other countries a PRI contains **23 B channels and one D channel**.

A **B channel** operates at **64 000 bits per second**. (56 000 bits per second in some parts of North America and other countries.)

A **D channel** operates at **64 000 bits per second** in a Primary Rate Interface. Note that this is different to the BRI where the D channel runs at 16 000 bits per second.

The animation below shows the multiplexing in a European Primary Rate Interface.

The Primary Rate Interface spends equal amounts of time transmitting data for each of the B channels and for the D channel because they all operate at the same speed.

Notice that there is one **timeslot** which has no channel assigned to it. This is reserved for your network supplier to use for diagnostic purposes. The D channel appears between B channel 15 and B channel 16.

Fractional PRI

ISDN service suppliers have the opportunity to supply interfaces where not all the channels are active. In most countries, when you take a Primary Rate Interface, you are charged a rental per channel. If you don't need all the channels that are available, you can ask to have some of these channels deactivated. This is known as **fractional Primary Rate**. The number of channels that you may request will vary according to the marketing policy of your service supplier.

What happens if you try to use more channels than your subscription permits? In ISDN, the network is the arbitrator of everything. When you want to place a call through an ISDN network, you send a request in your D channel to the network. The network will either attempt to satisfy your request or it will refuse it.

This user has only one B channel active in his subscription. He will be successful in making the first call. If he requests a second call while the first is still active, the network will reject the request.

Fractional BRI also exists, but is less common.

Protocols on B and D channels

The B channel is a neutral conduit for bits; so the meaning of the bits flowing in the channel must be understood by the users at each end.

If the two users are not using the same "language", they won't understand each other and so there is no meaningful communication. Imagine making a telephone call where the person who answers doesn't speak the same language as you. You can't even apologise for disturbing them!

The same is equally true of D channels. There are various "dialects" of the signalling protocol. You must be using the same dialect as your network in order to successfully communicate with it.

B and D channel protocols

You must use a protocol to establish meaningful communication across a channel. It is important that both parties to the communication use the same protocol.

This is particularly important for the D channel. Your signalling requests and responses must be understandable by the network. Even if your ISDN device and ISDN line are both functioning correctly, you might not be able to make successful calls if you're using a D channel protocol that isn't the same as the network's.

ISDN requires that you use a protocol defined by the ITU-T called **Q.931** for signalling in the D channel. However, there are several signalling protocols based on Q.931 in use round the world. For instance, NI-1 and 5ESS are used in North America while much of the rest of the world is now using EuroISDN (also called ETSI or DSS1).

You have a much greater choice of protocols for the B channel since the B channel is a neutral conduit for data of any type. You can use it to transmit any protocol you wish (eg. SNA or PPP). However, if the network doesn't understand the protocol it cannot give you any assistance if your call has to be delivered to a different type of network (eg. PSTN) where data conversion is required.

B channel charateristics

- It is important to remember that ISDN channels cannot be divided up into smaller units. Each is provided on an "all or nothing" basis.
- Two users communicating over a B channel have 64 000 bits per second available to them. There is nothing they can do to reduce this bandwidth.
- What about the situation where the two users find that 64 000 bits per second is not sufficient? The only solution is to add another B channel. This gives them 128 000 bits per second. They are **not** using a single B channel of 128 000 bits per second. (Don't forget that the speed of a B channel is defined as 64 000 bits per second. Anything which operates at a different speed isn't a B channel.)
- This means that they will have two parallel calls between them and the phone bill will show two simultaneous calls.

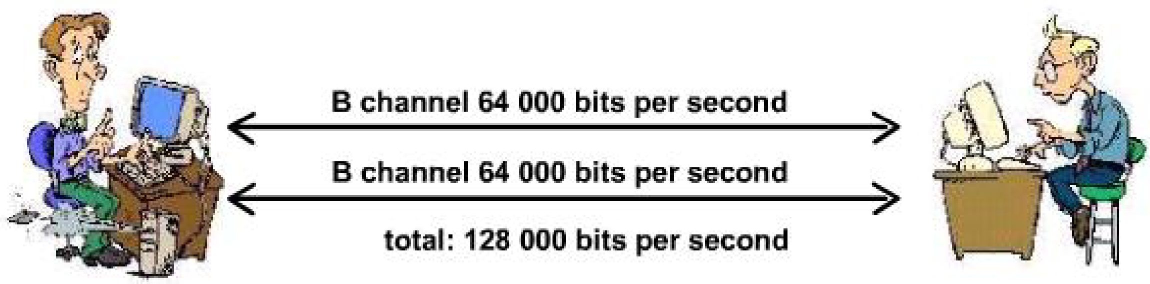

Using 2 B channels
Imagine that you're a user communicating with someone else, using two parallel B channels. Does the ISDN network care whether these two B channels are connecting the same two users or if they're connecting one user with two others?

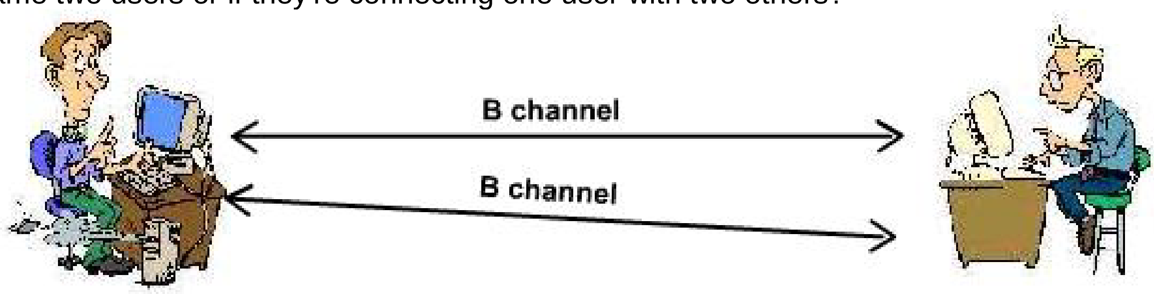

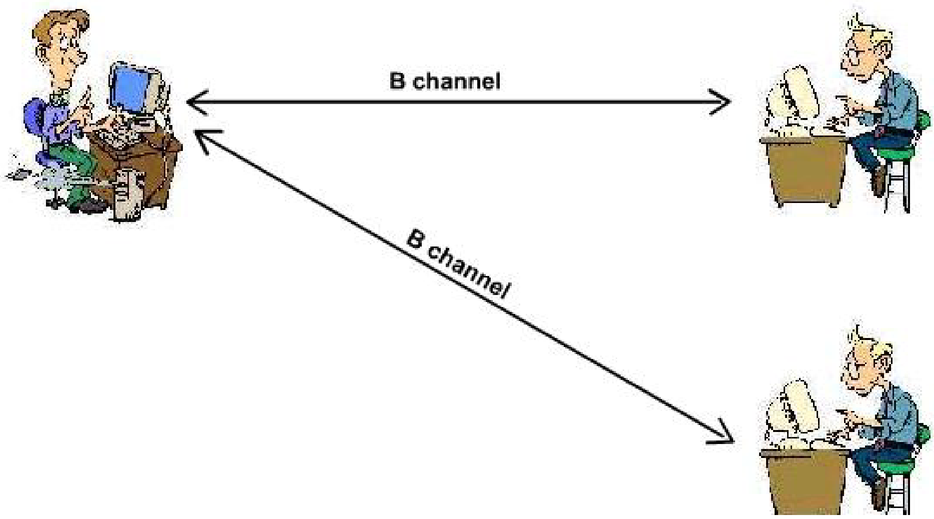

In the diagram below, the two users are connected using two B channels in parallel. The ISDN is able to route these B channels independently, because it takes no account of the fact the both channels connect the same pair of users.

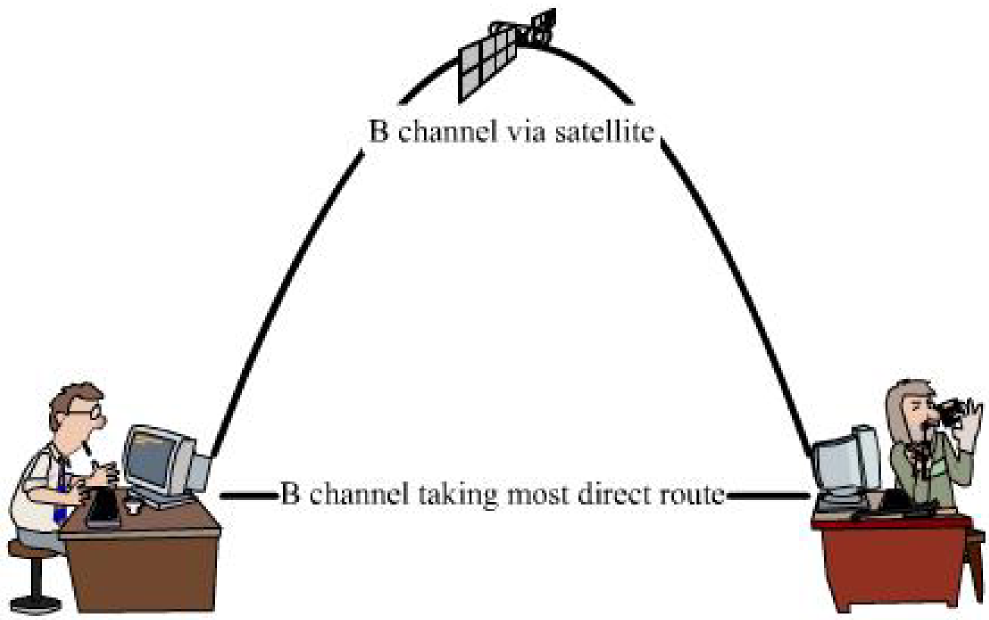

The speed of the two B channels is identical.

The time it takes for data to travel from one end of the channel to the other is, however, different.

One user transmits two items of data simultaneously. One is sent in the B channel which is routed via satellite; the other is sent in the B channel which takes the direct route. Will both items of data arrive at the same time?

The item of data which travels down the most direct path will arrive first. That which goes via satellite will arrive later because it has further to travel.

The ISDN makes no attempt to synchronise the data on the two B channels, possibly because it doesn't understand what the protocol in use. The two B channels are operating independently - the ISDN doesn't care that they're both connecting the same two users.

If the blue data item and the red data item together comprise a logical entity (eg. a picture on a web page), the receiving user cannot use the blue item until the red item has also arrived. It is his responsibility to understand this and to take appropriate action. Normally this will be done in software such as MLPPP.

Summary

- The D channel carries Network Services in the form of signalling. This is the way the user maintains his relationship with the network. Each user has one and only one D channel.
- The B channel carries Bearer Services which are the communication between two users. A single B channel cannot connect more than two users together.
- B channels and D channels share time on the interface.
- B channels cannot be sub-divided to provide less bandwidth
- More than one B channel can be used together to provide more bandwidth

Q/A

1. What is the D channel used for?

The D Channel is used for all the call maintenance traffic, such as call setup and call disconnect.

The D channel activity occurs between an ISDN user and the ISDN network. It does not pass across the network to the user at the other end.

The user at the other end of a call also uses a D Channel over which it conducts a separate dialogue with the network.

The network signaling is therefore separated from the B Channel which will carry the actual call itself. This connection does get established end-to-end between the two users.

2. Why can you not normally send data between users through the D channel?

A D channel connection is always between the user and the network, and not between the user and the user at the other end.

For this reason, using the D channel to transmit data is not normally possible since the connection effectively terminates at the network.

3. Why is there only one D channel in each interface?

There is only one D channel because that is all that is ever required.

The D channel only ever transmits network signaling information. Since, as far as ISDN is concerned, there is only ever one user and only ever one network, there is no need for more than one D channel.

4. Name two significant differences between B and D channels.

There is only ever one D channel per interface, whereas there can be multiple B channels

The D channel carries network signaling; the B channel carries the call itself

A D channel connection is always between the user and the network; a B channel connection will be between a user and any other user

5. What is the maximum number of simultaneous connections supported by a Basic Rate Interface?

Two.

There are two B channels within a Basic Rate interface, and both can be active at the any one time.

6. What is the maximum number of simultaneous connections supported by a Primary Rate Interface?

30 (or 23 in North America).

Primary Rate is supplied with either 30 (or 23) B channels and all of these can be active simultaneously.

7. If two ISDN channels are used together, each connecting the same pair of users, how many calls are being made?

Two. Even though the calling party and the called party are the same, there are two channels being used and two calls are being made.

This is important to remember for two reasons:

1) The setup for each call needs to be correct, and not having the details exactly right for the second channel can often be a reason for the failure of multi-channel calls

2) The advantage of the speed of a 128K link always need to be weighed against the cost of having TWO simultaneous calls being charged for.

ISDN Networks

In this section you will learn

- How ISDNs are connected together
- How users are connected to a public ISDN through an ISPBX
- How configuration errors in the ISPBX can prevent successful connections

ISDN and PBX

A normal PABX (Private Automatic Branch Exchange) handles only speech calls, whereas an Integrated Services Private Branch Exchange (ISPBX) is connected to the ISDN and itself provides ISDN extension lines.

The ISPBX has the ability to switch incoming calls from outside directly to its extensions. Being digital, it also has the ability to route not only speech calls but data, video, high quality audio, and Group 4 Fax. These calls can also be connected internally from extension to extension.

There are special characteristics about the relationship between individual users, the ISPBX, and the public ISDN. Understanding this relationship is essential when attempting to diagnose problems which can arise when using ISDN devices connected through an ISPBX instead of directly to an public ISDN line.

In the diagram below, how many users of the public ISDN network are there?

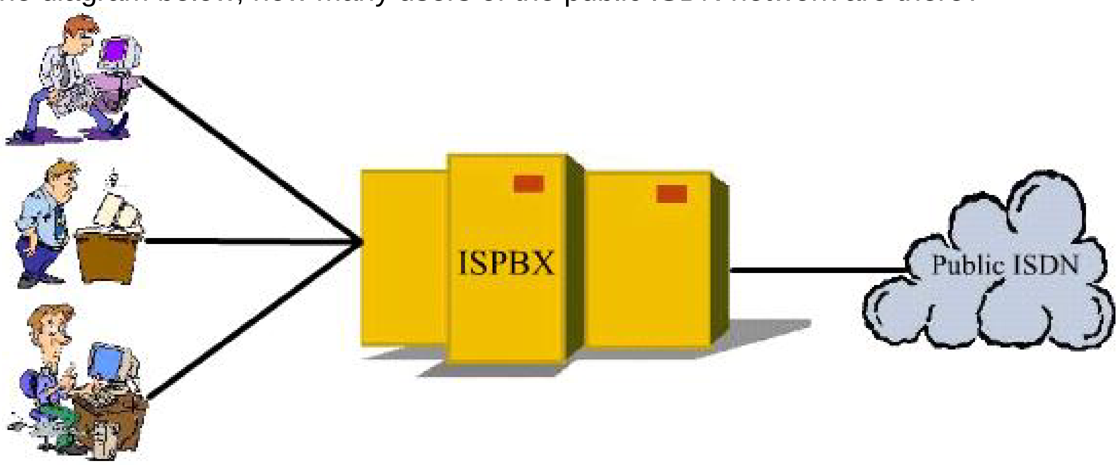

You can see that there is only one ISDN interface connecting the ISPBX to the public network. This means that there is only one D channel, which means that, as far as the public ISDN network is concerned, there is only one user.

"What about the real users?" you ask. "If they're connected with ISDN, what are they connected to?"

Assume that each of the three users shown has a Basic Rate Interface on the ISPBX. Each therfore has a D channel that terminates in the ISPBX. Therefore each is connected to an ISDN.

This network is **inside** the ISPBX.

The ISPBX acts as if it has **both** a network and a user inside.

The real (physical) users connect with the network inside the ISPBX over their D channels. If they need to communicate with the public ISDN network (for example, to ring home), the "virtual" user inside the ISPBX sends the service request to the public network on their behalf.

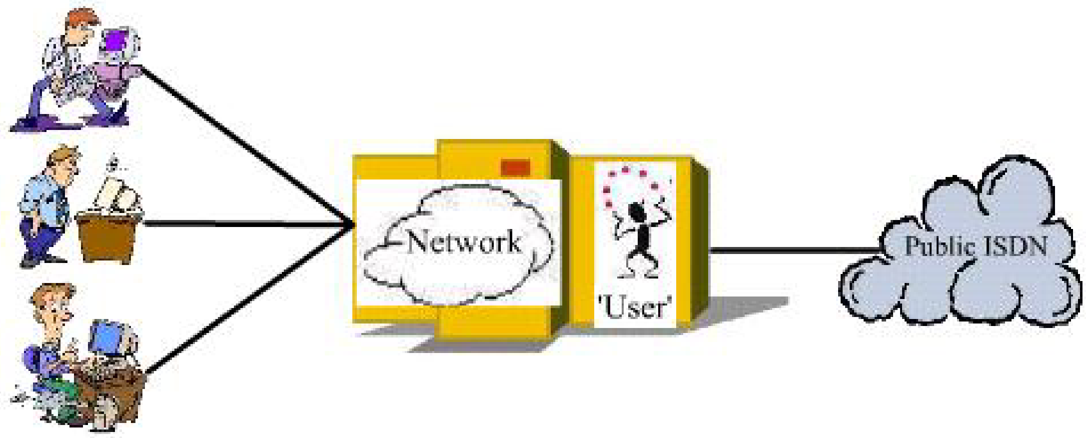

The services provided by the ISPBX are often different from those provided by the public network. This can cause problems!

The network service request that a user sends on his D channel arrive in the ISPBX's internal ISDN. This is **not** passed directly to the public ISDN. The ISPBX interprets the request and decides what to send on the D channel that connects it to the public ISDN network.

Summary

- All ISDN users need to be connected to an ISDN network
- This network can be the public ISDN, or an office-based digital PABX (an ISPBX)
- An ISPBX needs to be connected to the public ISDN
- The various public ISDNs are connected together
- A user's call request might be modified by any of the intervening networks

Q/A

1. When multiple users on an ISPBX are making calls, how many users does the public ISDN network see?

One. The calls to the network actually come from the Primary Rate interface on the ISPBX. As far as the network is concerned, the ISPBX is a single user with multiple channels available with which to make multiple calls.

2. Do users on an ISPBX have a D channel connection to the ISPBX or to the public ISDN network?

The users on an ISPBX do not communicate directly with the public network at all.

The users see the ISPBX as the network, and each active user has a D channel connection with the ISPBX, not the network.

3. An ISPBX can behave like a single user and a network - true or false?

True.

It depends which interface on the ISPBX is being described.

To the users connected to an ISPBX, it appears as a network, and they each can establish D channel communications with it.

As far as the interface to the public network is concerned, however, the network simply sees the ISPBX as single user with multiple channels available.

4. A user on an ISPBX can use different D-Channel protocols to those in use by the public ISDN network connected to the ISPBX - true or false?

True.

Since the ISPBX is effectively a private network to which all the users are connected, the ISPBX is in total control of the ISDN protocol it wishes to use.

The information from the users gets presented to the public network by the ISPBX itself. The ISPBX therefore needs to communicate externally to the public network in whatever protocol is being used by the local ISDN provider. Internally, however, it can, and frequently does, deploy ISDN using a proprietary implementation of the protocol.

This is usually in order to add special features for use within the private network. However, this also means that certain functions which standards-based ISDN equipment expects to be present may not be implemented in the normal way on the ISPBX.

These differences can frequently be the cause of problems arising when using ISDN equipment via an ISPBX. The best method of diagnosing these type of faults is to get access to a standard public ISDN line and see if the problem persists. If it doesn't, as is likely, then the ISPBX will need to be re-configured or upgraded.

Bearer Capabilities

In this section you will learn:

- Bearer Capabilities in depth
- The significance of Bearer Capabilities in diagnosing problems
- A practical example of using Bearer Capabilities to resolve a problem

Bearer Capabilities - What are they?

Although the ISDN doesn't always need to know what protocol is in use in the B channel, there are circumstances where this information is useful.

For instance, if you make a telephone call over ISDN, it is useful to tell the network that this is a voice call so that it can connect your call with an analogue telephone in the PSTN. If the ISDN network thinks that your B channel contains a protocol unknown to it, then it can only connect your call directly to another ISDN line.

You have the opportunity to tell the ISDN about the protocol that you're using in the B channel when you request the call. This information is added to the signalling information sent on the D channel when the call is requested. This call information is often called Bearer Capabilities.

The receiver of the call can also see the bearer capabilities when a call is offered to him by the network.

Generally, there is no possibility to negotiate bearer capabilities. You must decide what bearer capabilities you want to use before placing the call. Bearer capabilities are fixed for the lifetime of the call and cannot be dynamically changed. If you attempt to place a call without specifying the bearer capabilities, the request will be rejected by the ISDN.

Using Bearer Capabilities

Understanding Bearer Capabilities and their implications becomes significant when diagnosing problems, particularly when using diagnostic traces.

An example of the problems that can occur involves fax calls not being received correctly because there are two possible types of Bearer Capability for a fax call.

- 3.1kHz analogue call
  This often happens where the call originated inside the PSTN. The ISDN has no way of knowing exactly what type of equipment (telephone, fax, modem, etc.) placed the call, so it uses this "catch-all" bearer capability.
- Fax Group 3 call
  This is often used where the originates with a fax machine that is directly connected to an ISDN.

For the call to be successful, the chosen Bearer Capabilities must be supported by all equipment throughout the path taken by this call across the network. If any single piece of equipment does not support the requested Bearer Capabilities, the call will not be connected. In general, there is no support for negotiating Bearer Capabilities.

We therefore need to ask the questions:

- Do all the ISDN networks traversed by the call support the requested Bearer Capabilities?
- Does the PBX equipment at each end support these Bearer Capabilities?
- Has it been enabled?
- Has it been enabled for the port in question?

Bearer Capabilities - An example – 1

Consider this example of the subtleties of Bearer Capability requirements.

A user of an ISDN PBX (ISPBX) may have no problem making analogue calls, yet experience difficulties sending faxes. The Bearer Capability required to send a fax might be:

- Enabled for the port in use on the remote PBX
- Available as a facility on the local and remote PBX
- Supported by the local and remote providers
- Enabled for various ports on the local PBX

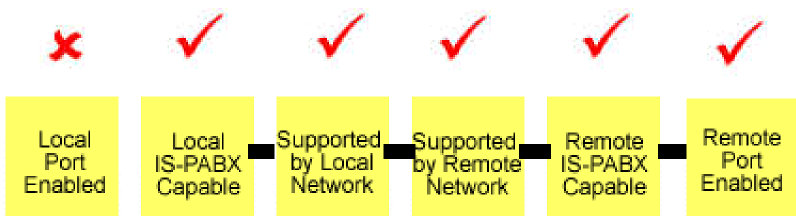

- But **not** enabled on the particular PBX port experiencing the problem.

The chain of enabled Bearer Capabilities must be complete for the service to be operable. In this instance, the PBX will refuse to recognise the outgoing fax calls from the local port.

Bearer Capabilities - An example – 2

Examining a diagnostic trace for this scenario would reveal this, since the cause code reported would indicate that the service being requested is "not available" or "not permitted". It is potentially available, of course, but not for this call on this port.

The solution in this instance would be to enable the particular function required on the port in question. This would normally require that the PBX be re-configured:

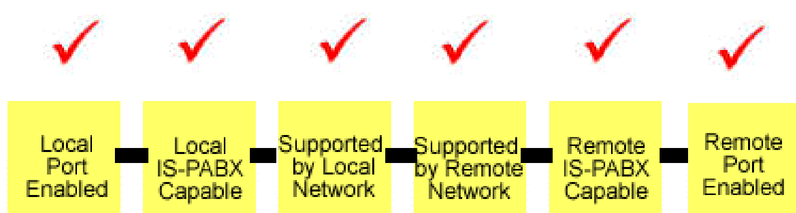

Bearer Capabilities - An example – 3

PBX manufacturers often extend or amend the rules for internal ISDN communications, so the cause code returned may not be exactly the same as the international standard.

The point to note in the example given, is that neither the ISDN equipment, nor the public network, nor the intended recipient, is the cause of the failed communication. Sending the same fax call via a **direct** ISDN line would prove this.

Bear in mind also that the whole scenario could be repeated at the remote end of the connection, with the supplier, the PBX, and the remote port, all needing to be aware of the Bearer Capability being employed.

Summary

- The Bearer Capabilities required for a call need to be available **throughout** the path of a call for the call to succeed.
- Diagnostic traces reveal Bearer Capabilities to assist problem identification.

Q/A

1. A fax call call might be designated as a generic analogue call or specifically as a fax call - true or false?

True.

A call originating from a fax machine directly connected to the ISDN network is likely to have 'Group 3 Fax' as its call type.

A fax call originating on the PSTN network will not have any bearer capabilities to start with, and will get allocated an automatic category of call type called 'Generic Analogue' by the ISDN network when the call is transferred from the PSTN.

2. How can you find what Bearer Capabilities are specified on an incoming or outgoing call?

Diagnostic traces which send the progress of a call setup and disconnect to a readable log will reveal the Bearer Capabilities of a particular call.

Eicon provide a utility with all their DIVA ISDN products called DiTrace which facilitates this process for diagnosing problems.

Voice, Fax, and Modems

In this section you will learn about

- How ISDN and analogue networks interact
- Digital and analogue conversions in the network
- Where and how these conversions take place
- MODEMs and CODECs
- How they impact making fax and modem calls
- ISDN equipment types and their fax and modem capabilities

Analogue Calls and ISDN

The key characteristic of ISDN is that it is a digital network. However, many of the devices and networks with which an ISDN user needs to communicate are not digital but analogue. In order for these two types of device to communicate, the information that they are exchanging must be converted from one form to the other.

In fact, except for data calls between computers to across the ISDN network, almost all other types of calls - voice, fax, modems - will all involve some kind of conversion from digital to analogue, or vice versa.

Much of this conversion takes place without the user's knowledge or intervention and is handled by the networks and devices involved. However, there are instances where an understanding of what is involved will assist in making successful connections and diagnosing problem areas.

You need to pay careful attention to the requirements of the ISDN device in use, particularly when sending and receiving faxes. This section provides the background to the various scenarios involved, and the practical implications for the different types of ISDN device that are available.

Voice over ISDN

ISDN is a **Digital** network. Everything (including sounds such as voice and modem signals) is carried as a stream of bits.

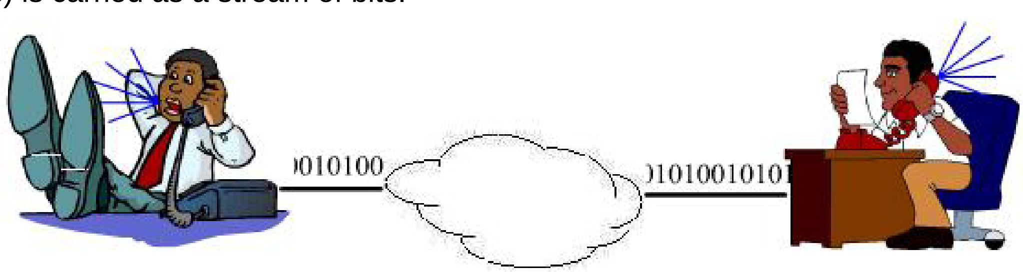

This means that ISDN telephones need to be able to digitise and "un-digitise" sounds. This is performed by a device called a **CODEC** (Coder-Decoder) which is located inside the telephone. The CODEC translates the sounds into bits in one direction, and translates bits into sounds in the opposite direction.

The analogue signal originating in the microphone of the telephone handset is sampled and transformed into a stream of bits (64 000 of them every second) that is placed on the B channel.

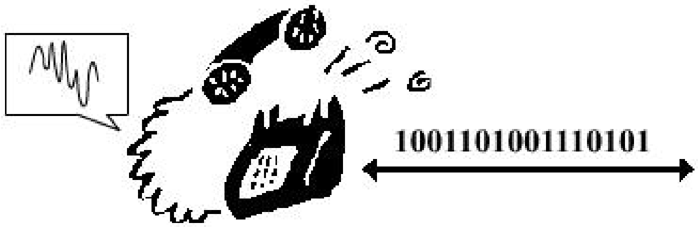

Similarly, the incoming bit stream from the B channel is converted back into an analogue signal and sent to the ear-piece of the handset.

A B channel is **full duplex**, which means that it can carry data in both directions at once.

The ability to make voice calls from one ISDN telephone to another over a digital B channel is indeed useful, however, the majority of telephones currently installed worldwide are analogue devices which are not connected to an ISDN.

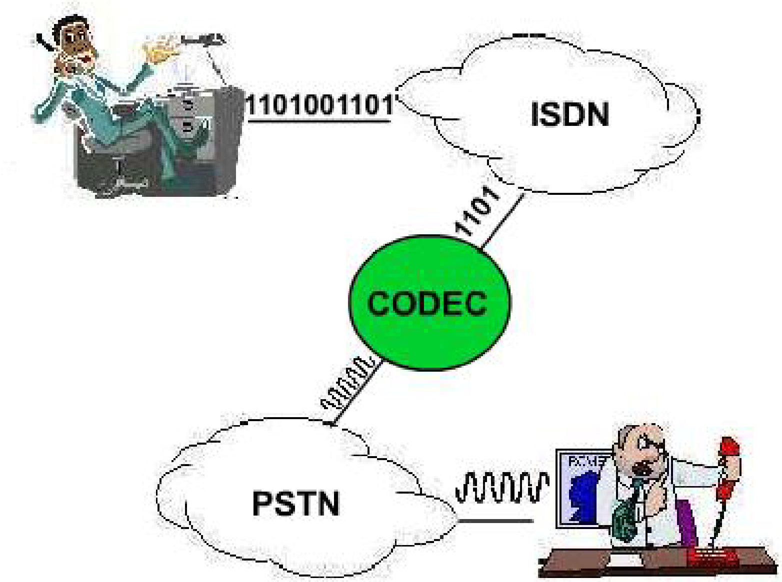

Fortunately, you can make calls between the two networks. For this to work successfully, there has to be a conversion between the bit stream in the B channel and the analogue signal required by the PSTN.

CODECs are located at the boundaries of the digital and analogue networks.

Fortunately, you can make calls between the two networks. For this to work

The CODECs inside the network and the telephone **must** use the same rules when formatting the bit stream that represents the users' voices. Provided both devices doing the same processing, then the information can be converted by applying the same rules in reverse.

Given the presence of the CODEC in the network, and adherence to the correct protocols, any device that can be used on the PSTN, such as a modem or a fax machine, can also pass calls into the ISDN.

Analogue fax and modem over ISDN

Another important idea is introduced here; this is the Terminal Adapter.

A Terminal Adapter (TA) is always necessary to connect non-ISDN devices (such as a serial port of a PC) to the ISDN. However, a TA can also contain a CODEC if it is intended to support analogue phones, fax machines and modems.

In diagram below, the modem at the top left can plug in to the POTS ports on the TA. The TA will then convert sounds generated by the modem on its POTS port into a bit stream (and vice-versa). This bit stream is identical to that created by an ISDN telephone; that's to say it represents **sounds**.

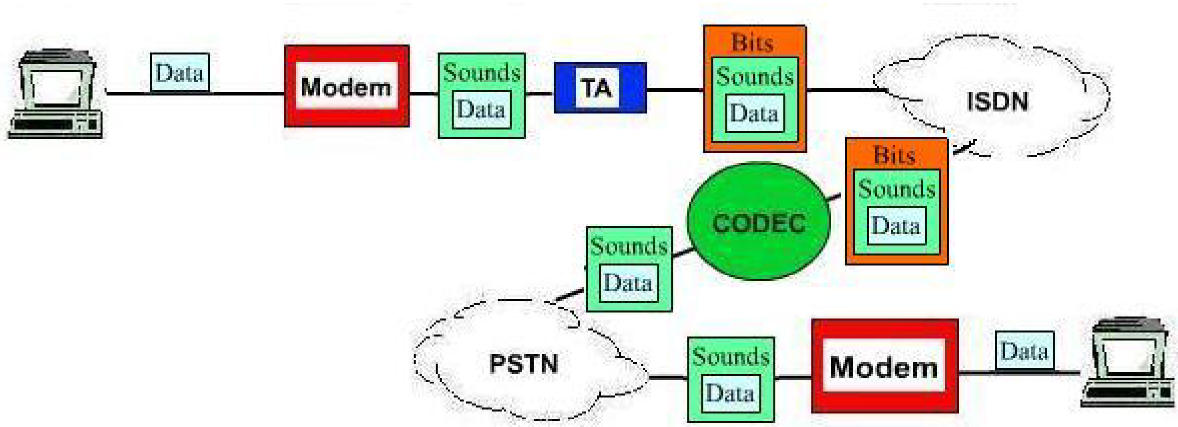

Starting in the bottom right-hand corner, data leaves the PC as bits that are converted into sounds by the modem. We now have data encapsulated in sounds.

These sounds cross the PSTN network until they are encapsulated inside a bit stream by the CODEC at the boundary between the ISDN and the PSTN. This bit stream is then passed from the ISDN network to the Terminal Adapter, which contains a CODEC that converts the bit stream back into sounds. These sounds are sent to the modem at the top left, which converts this back into the original data that entered the modem at the bottom left.

Starting in the bottom right-hand corner, the data leaving the PC is converted into sounds by the modem. We now have data encapsulated in sounds. These sounds cross the PSTN network until they are encapsulated inside a bit stream by the CODEC at the boundary between the ISDN and the PSTN. This bit stream then passes through the ISDN network to the Terminal Adapter, which contains a CODEC that converts the bit stream back into sounds. These sounds are sent to the modem at the top left, that converts this back into the original data.

The process runs in the opposite direction to send data from the PC in the top left-hand corner to the PC in the bottom right-hand corner. This appears to be a lot of work: the data sent across the ISDN has been encapsulated twice.

Double Encapsulation of Analogue Modem Data Across an ISDN:

Between the PSTN and the ISDN there is a CODEC that samples the sounds the modem is making and creates a bit stream representing these sounds. This means that the data that originally entered the modem is now encapsulated as bits representing sounds that represent the data.

Why couldn't the CODEC, merely interpret the sounds and recover the original data - like this?

This would appear to present no problem for the ISDN which is digital and can therefore carry the data passed into the modem (which is digital as well).

The next step in evolving this configuration is to use an internal modem in the PC.

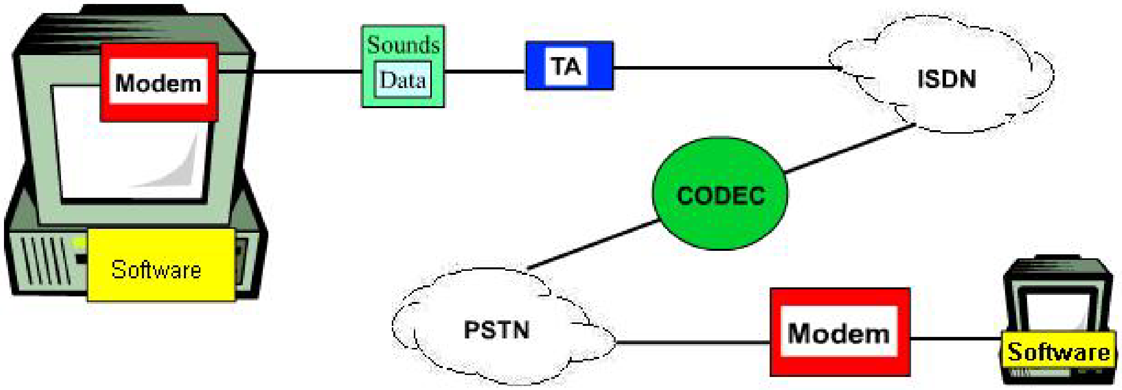

Logically this is no different from the previous scenario.

We could replace the internal modem with an internal ISDN adapter and a driver which appears to the application software to be a modem. In reality, this driver combines the functions of both the modem and the CODEC. This driver is known as a **soft modem**.

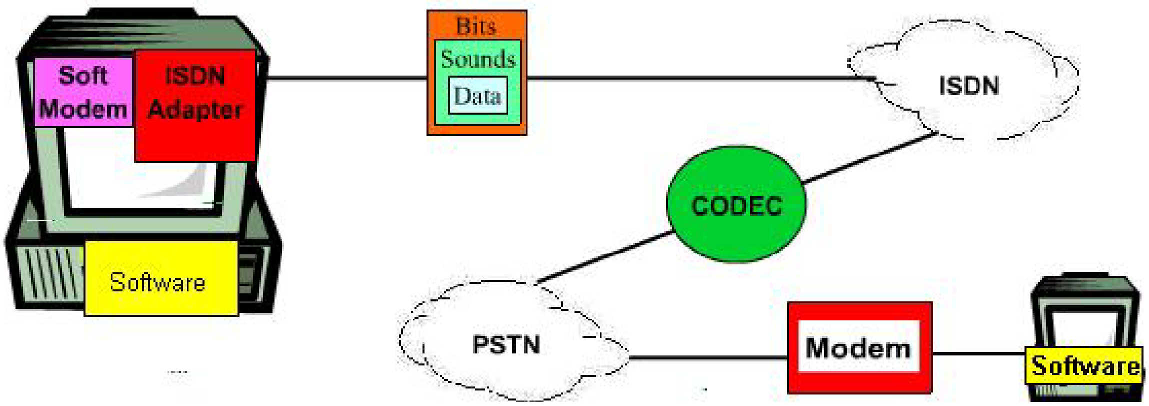

A soft modem requires a lot of processing power, since it has to operate in real time. The PC is therefore likely to appear a bit sluggish while the connection is active. The advantage of this solution is that it can be used with an inexpensive passive ISDN adapter.

The final step in evolving this scenario is to use an ISDN adapter with has a DSP (Digital Signal Processor). This takes the very heavy processing load due to the soft modem away from the PC. ISDN adapters with a DSP a generally more expensive than passive adapters.

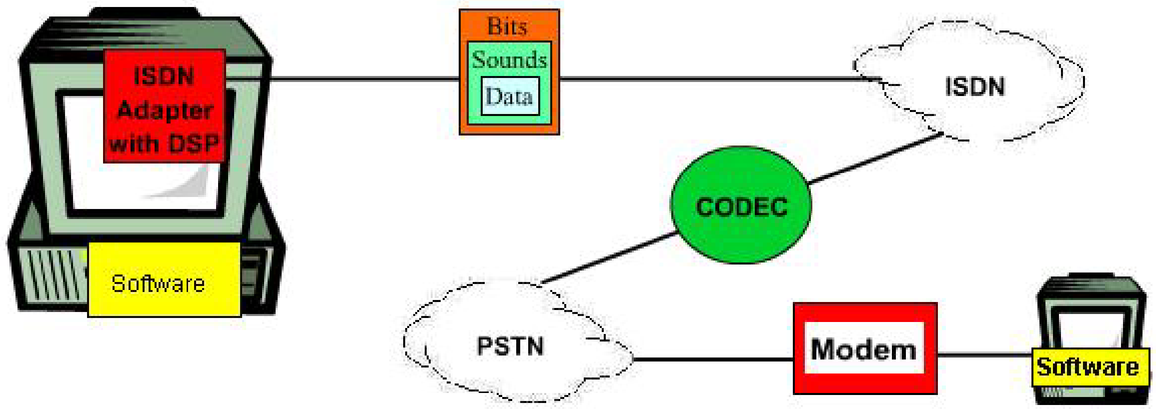

Soft-fax and Soft-modem

This means that an ISDN adapter with an on-board DSP can also communicate with any device that contains a modem. In the diagram below, you can see that we've added a fax machine.

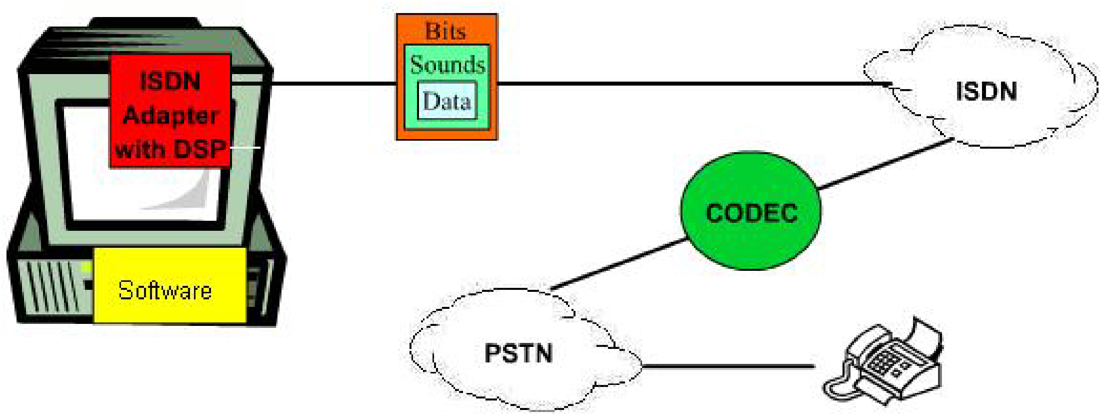

To perform the job of a modem and a CODEC at the same requires a large amount of processing power. DSPs are very powerful processors. Nevertheless, you need one DSP for each B channel for which you want to use this technique. There is, however, no reason why you couldn't use an ISDN adapter that has a single DSP and a soft modem driver to handle two modem calls at a time.

Summary

- In ISDN networks everything is carried as a stream of bits.
- Converting digital telephone signals into voice, and vice versa, is done by CODECs.
- Converting data into analogue telephone signals, and vice versa, is done by modems.
- A CODEC is the device that allows telephony between an ISDN and the analogue network.
- This same CODEC can be used to allow modem and fax calls to cross the same boundary.

- Modems and CODECs always work in pairs - they can be nested together but each must have a partner.
- A Terminal Adapter is an interface between the ISDN and any non-ISDN device, such as a computer or an analogue phone
- If analogue devices need to be connected to the ISDN, then the Terminal Adapter will need to perform the function of a CODEC, and have analogue (POTS) ports available for modems and fax machines to plug into.
- Internal ISDN adapters can use dedicated chips on the card, or software running on the PC, to implement the digital and analogue conversions, thereby removing the requirement for any physical analogue devices like fax machines and modems.
- There are three kinds of device used in data communications with ISDN:
    - Conventional Modems - these require Terminal Adapter to connect to the ISDN
    - ISDN Modems - which combine the functionality of CODEC and modem
    - Terminal Adapter - allow analogue devices to connect to the ISDN and contain a CODEC for this purpose
    - ISDN Adapters - merely pass a stream of bits between a protocol driver and the ISDN

Q/A

1. What is the device inside an ISDN telephone that converts voice into digits?

A **CODEC.**

A word deriving from the two functions performed : **CO**der - **DEC**oder

2. What happens to these bits when they are received by the ISDN telephone at the other end?

The CODEC within the ISDN telephone at the other end will translate the incoming digital signal into an analogue format.

It can then be reproduced as voice through the speaker in the earpiece.

3. What device converts the signals between the ISDN and the PSTN networks?

This again is a **CODEC.**

It is operated by the local ISDN provider, and will be situated at the boundary of the ISDN network and the PSTN network.

This device enables calls to pass from one network to the other by converting bits to waves, and vice versa.

4. For all the above conversions to succeed, what must be common?

CODECs are not simply generic devices. They need to perform their conversion according to a standard so that a complementary **CODEC** at the receiving end can apply the same standard when it does the reverse conversion.

The **CODEC** will encode the data according to a fixed set of rules - as long as the receiving **CODEC** applies the same set of rules in reverse, then the data can be reconstituted exactly as it was originally transmitted.

5. What is the purpose of an ISDN Terminal Adapter?

The purpose of a Terminal Adapter is to connect items of equipment that are not ISDN-aware to the ISDN network.

An ISDN telephone, for example, is ISDN-aware, and can be connected directly to the ISDN network.

The serial port on a PC, on the other hand, is not ISDN-aware, and will therefore require a Terminal Adapter to provide the appropriate interface between the PC and the ISDN network.

6. Does a Terminal Adapter always contain a Codec or Modem?

A Terminal Adapter does not necessarily need to have a CODEC inside it.

For example, a basic Terminal Adapter could connect a PC via its serial port to an ISDN network without any requirement for a CODEC.

The bit stream from the serial port is digital, as is the bit stream that goes out to the ISDN. Although a protocol conversion does indeed occur, it is not an analogue to digital conversion, so no CODEC is required.

On the other hand, when an analogue device, such as a fax machine, needs to be connected, then an analogue to digital conversion is necessary and a CODEC will be required.

So, simple Terminal Adapters do not need to have CODECs, but they cannot connect analogue devices to the ISDN.

More sophisticated Terminal Adapters that are able to connect analogue devices to the ISDN do need to have a CODEC.

7. What are the two options for implementing fax using an internal ISDN card?

To send and receive faxes, an ISDN card needs to perform digital conversions, as if it were an analogue modem.

It can do this by having an on-board Digital Signal Processor perform the conversion, or it can off-load the task to the processor of the host PC and use a 'soft modem' application.

The trade-off between the two is that DSP chips are sophisticated technology and are therefore expensive, whereas using software to perform the conversion can take up a significant amount of the processing power of the computer.

Terminology

In this section you will learn:

- The naming conventions used within ISDN
- The essential difference between the key items
- The significance of the various reference points
- The practical implications of international variations

Where the network terminates: The NT1

Most ISDN services are provided through a device known as an NT1. This stands for Network Termination type 1. The device is provided by your network operator and is part of the ISDN. This is the point at which you access the network The upper picture on the right shows the NT1 supplied by British Telecom.

In North America, the ISDN service provider does not provide the NT1 for you. You are provided with a simple socket to which you may attach a single ISDN device. If you wish to attach more devices, you must purchase an NT1 and connect it to this socket.

In most of the rest of the world, the NT1 is provided by the ISDN service provider and normally has two sockets into which you can plug ISDN devices.

The NT1 is an active device, with quite complex electronics that handle the transmission of the 144kbps data stream to and from the ISDN service provider's switch or exchange.

An NT1 cannot make calls by itself. You need to plug in Terminal Equipment (TE) or a Terminal Adapter (TA) to do anything useful.

Some NT1s (sometimes called NT1+ or Super-NT1) also have analogue ports on them so you can plug in an ordinary analogue phone. The lower picture on the right is the British Telecom Highway Super-NT1. The blue sockets provide ISDN; the white sockets are for connecting analogue devices.

Types of equipment that can be connected to ISDN: TE & TA

Terminal Equipment Type 1 (TE1): understands how to interface to the ISDN network

- Can connect directly to the ISDN
- E.g.: ISDN interface card or USB device

Terminal Equipment Type 2 (TE2): understands nothing about ISDN

- Must use a Terminal Adapter (TA) in order to be connected to ISDN.
- Examples: modem, fax, analogue telephone, PC Serial Port. (Perhaps even a television camera).

A **Terminal Adapter** is required to connect a TE2 device to the ISDN network.
A Terminal Adapter is by definition the equipment needed to connect a TE2 to an ISDN.

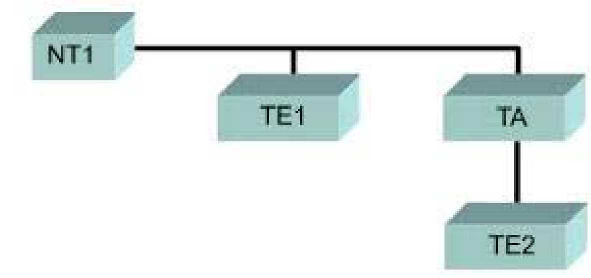

The NT1 Plus or Super NT1

The 'NT1 Plus' or 'Super NT1' provides additional services to the user by incorporating one or more Terminal Adapters (TAs) in the same enclosure as the NT1.
Examples:

- A TA built in to the NT1 housing may provide one or two analogue phone ports so that you can plug in an analogue telephone, modem, or fax machine.
- A serial or USB port provided by another type of TA allows you to connect your PC via a serial cable.

ISDN service providers have many local names for ISDN services that include these capabilities within the same enclosure as the NT1

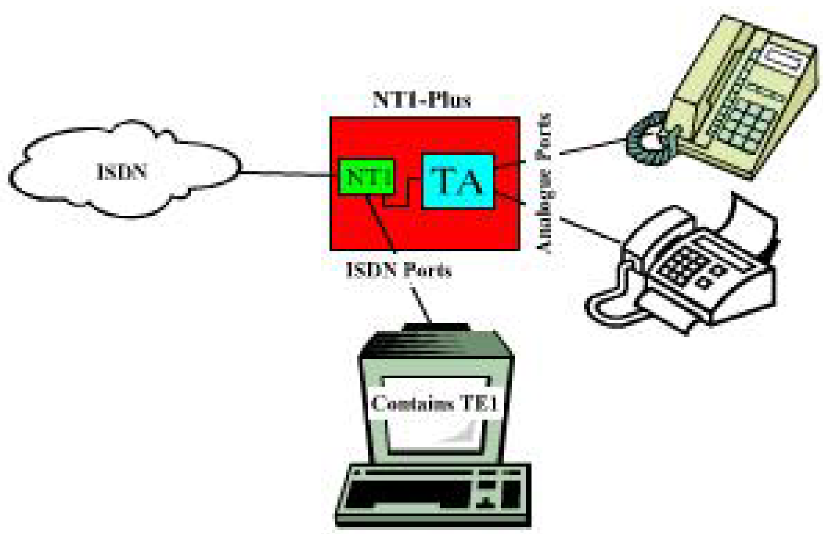

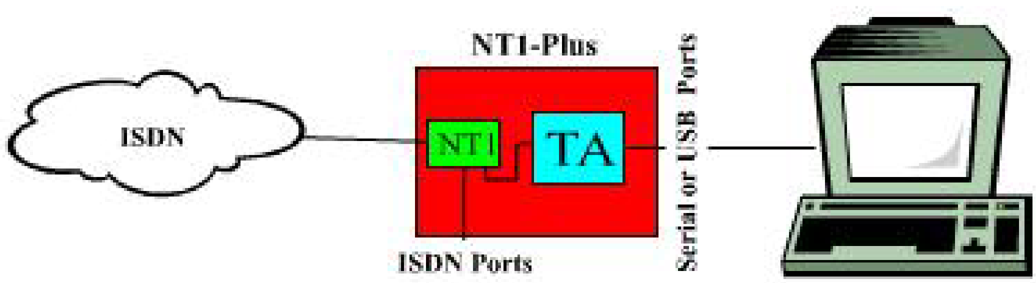

The location of the S/T interface

ISDN standards define certain key interfaces of the network as Reference Points.

Reference Point T is on the user's side of the network termination and allows a single TE or TA to be connected.

Reference Point S allows more than one TE or TA to be connected. It only exists for Basic Rate ISDN.

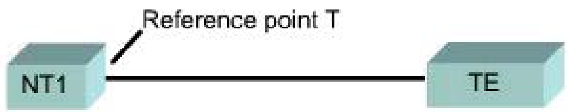

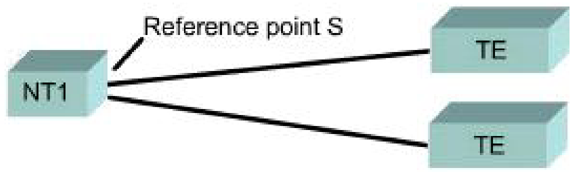

In this configuration, Reference Point S and Reference Point T are effectively in the same place. The only distinction is the number of devices attached and so this reference point is often referred to as S/T.

The S-bus

The passive bus illustrated below is an example of Reference Point S. For this reason, it is often known as an S-Bus. This is the normal configuration for EuroISDN. Up to eight devices can be connected to the S- Bus, any two of which can be be active simultaneously.

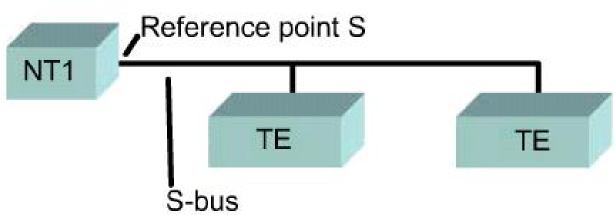

The location of the U interface

In North America (and some other countries), where the NT1 is not provided by the ISDN service provider, the Reference Point U is used to define the end of the wires provided to the user by the service provider. There is no international standard that defines Reference Point U.

Only one Terminal Equipment may be attached to Reference Point U. The TE must be designed to operate on Reference Point U.

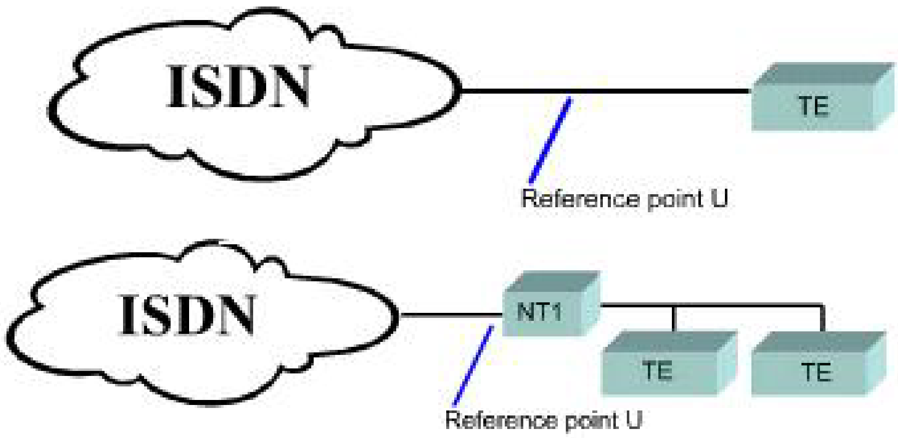

It is therefore necessary for the user's own NT1 to be connected to reference point U. This will be necessary if the user wishes to attach more than one device. He would have to use equipment designed to operate on the S or T Reference Points.

Reference Point U is explicitly not defined in EuroISDN, because the NT1 is always provided and permanently connected to the wires by the service provider.

Summary

In this section you have learned that:

- ISDN equipment is called Terminal Equipment
- There are two classes of TE: TE1 and TE2.
- TE1 is ISDN aware equipment (e.g. an ISDN telephone)
- TE2 equipment is not ISDN aware (e.g. an analogue telephone)
- A Terminal Adapter is the equipment required to connect TE2's to an ISDN network
- An NT1 is the termination of the network
- Reference Point S/T is on the user side of the NT1
- Reference Point U is the network side of the NT1
- In EuroISDN, the user connects to Reference Point S/T and the network provider supplies the NT1
- In North America, the user connects to Reference Point U and has to provide the NT1

Q/A

1. What are the two types of equipment that can be used on ISDN called?

The two types of equipment which can be connected to ISDN are called TE1 and TE2.

The TE stands for Terminal Equipment.

2. What is the essential difference between them?

Type TE1 is equipment that can be connected directly to the ISDN. For example an ISDN telephone.

Type TE2 is equipment that cannot be connected directly to the ISDN and requires a Terminal Adapter. For example a PC.

3. What item of equipment would be required to connect a modem to ISDN?

A terminal adapter.

4. What is the name of the reference point that is located on the user side of an NT1?
The reference point on the user side of the NT1 is called Reference Point S/T

5. Does the network provider supply the NT1 in North America?
No, in North America the user is required to supply the NT1.

6. In order to be able to plug analogue devices into sockets on an NT1, what extra item of equipment would need to be included within the the NT1?
The NT1 would need to incorporate a Terminal Adapter.

7. What are these types of NT1 known as?
An NT1 which includes a Terminal Adapter is called a 'Super NT1' or 'NT1-Plus'

Device Addressing

In this section you will learn:

- Why a device needs a unique address
- What the address is used for
- Forms of address management
- The difference in address management between PRI and BRI lines
- Practical implications of mixing address types

Why does ISDN need addresses?

A Basic Rate Interface (BRI) line may have as many as eight devices attached.

All these devices may use the common D channel at the same time. For example, one device may be trying to make a call at the same time as a call is being received and answered by another device.

The ISDN network therefore needs to be able to conduct 'conversations' with each device independently.

To do this successfully, each device requires its own address, called a **Terminal Endpoint Identifier** or TEI. <u>This address is **not** the telephone number</u>. Telephone numbers belong to the line, not to any device attached to the line.

Fixed and dynamic addressing

A PRI line can have only one device connected. This device will have a single, predefined address which, in practice, is always zero. When a ISPBX is attached to the PRI line, it will perform its own call management and call routing functions for the devices attached to it.

A BRI line can have up to eight devices connected to it. Device addresses may be pre-configured or dynamically allocated.

- In dynamic addressing, each device has to request an address from the network before it can perform any signalling activity.
- With fixed addressing, the address of each device has to be configured to match the addresses pre-configured for this line.

ISDN service suppliers do not allow both fixed and dynamic addressing on the same line. In practice, an ISDN line where the device addresses are fixed will only permit a single device. The is no convention to describe this type of line configuration, but ISDN service providers sometimes give this a name that reflects this type of configuration.

Dynamic address assignment

Dynamic address assignment means that the device must ask the network for an address before it can perform any signalling.

The Q.921 standard (also called LAPD) that defines the link level protocol used on the D-channel incorporates a mechanism by which the device can request an address.

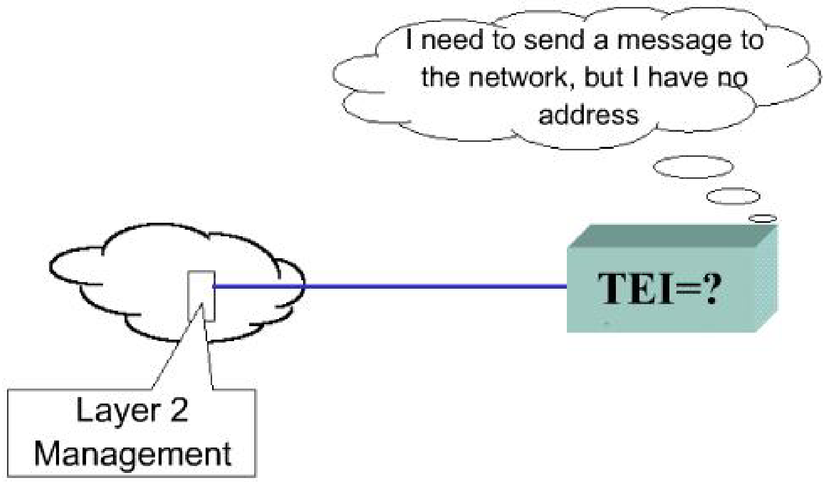

In this case, the device which currently has no address is able to request one from the network. The network responds with an address for the device which it will use for all future communications.

TEI assignment - potential problems

Problems will arise if the line is configured to expect a fixed address device, and the device connected is expecting to be allocated an address.

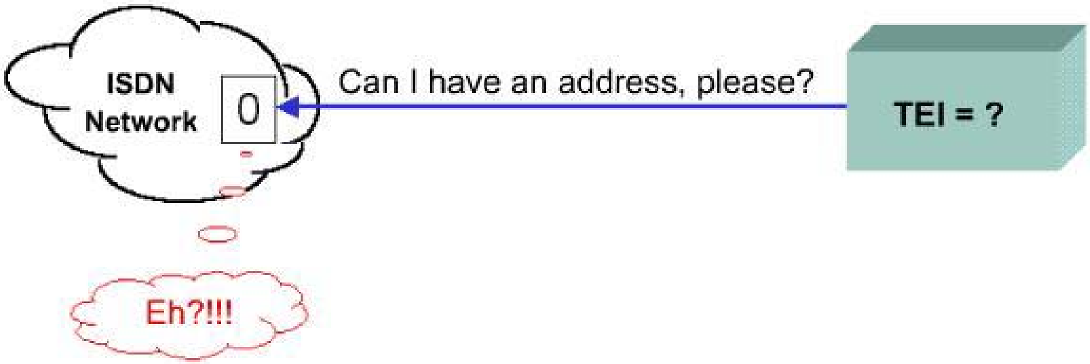

In this example, the network will attempt to activate device 0, which doesn't exist. The TE, which currently has no address, ignores this activation request because the address doesn't match his own. It waits for the allocation of a dynamic address. The net result of this is that the communication between the network and device is never initialised.

Broadcast address

In addition to fixed and dynamic addresses, there is also a universal address. This broadcast address is used when a message has to be heard by all devices sharing a given D channel.

Messages sent by the network to the broadcast address are heard even by devices currently without an assigned TEI.

Broadcasts are mainly used by the network to signal an incoming call - that way, any device connected to the line can take appropriate action - an ISDN telephone can ring, a TA could make attached analogue phones ring, or a device could decide to answer the call.

However, before a device without a TEI can answer a broadcast (to answer an incoming call, for instance), it must request and receive an address in the normal way.

Summary

In this section you learned:

- ISDN devices need to have unique addresses
- These addresses are known as TEIs
- An address is nothing to do with the telephone number
- PRI lines permit only a single device at address TEI 0
- BRI lines can have either fixed or dynamic address assignment
- Fixed address lines permit only one device
- Dynamically assigned lines can have up to eight TEIs
- Mixing TEI assignment types can result in communication failures

Q/A

1. How many devices can be connected to a standard BRI line?
Eight devices may be connected to a standard ISDN Basic Rate Interface.

2. What are the two ISDN device addressing methods?
The two types of device addressing in ISDN are known as:

Fixed Addressing and **Dynamic Addressing**

3. A device is configured for automatic TEI. What is its address immediately after being switched on - before it starts communicating with the network?
At this point, the device has no address - it will make a request to the network and the address will be dynamically allocated.

4. In practice, what is the maximum number of devices that may be connected to a BRI which expects a fixed TEI?
Only **one** device may be connected to a fixed TEI address ISDN line.

5. What is the significance of the TEI in B channel communication?
If the device is configured for one type of addressing and the network is expecting another type of addressing, then successful B channel connections will not be possible because the device will never initialise.

6. What address does an ISDN PRI line always use?
A Primary Rate line will always have the fixed address of **zero**.

7. Why does there need to be a broadcast TEI?
A broadcast address is necessary when the network needs to communicate with all the devices connected to an ISDN line.

The typical example of the need for this would be to signal to all devices that there is an incoming call to be answered.

Call Setup

In this section you will:

- See the step-by-step process for establishing and disconnecting an ISDN call
- See the exchange of messages between Terminal Equipment and the Network on both sides of a call.
- Understand how several devices attempting to answer a call are managed

Network Services

The interactions between the user and the network are all carried by messages in the D channel. These can be categorised into various groups

- Call Establishment
- Call Clearing
- Call Information
- Miscellaneous

The user at one or other end of the potential call negotiates with the network via his D Channel.

A successful call set-up between the two users will normally establish an end-to-end connection across the network on a B Channel.

Upon completion of the call, the disconnection procedure will also take place on the D Channel.

An understanding of the steps involved are essential for analysing problems encountered during call establishment.

Making a call

The following sequence of diagrams shows the process when a user requests a call through an ISDN network.

- The SETUP is a request to create a connection by the user on the left to a user on the right. The SETUP message usually includes at least the following information:
    - The called number (the destination number)
    - The call type (Bearer Capabilities)
    - It may also include other information, such as the calling number (the origin of the call).
- The network chooses a B channel for this half of the B channel and uses the SETUP_ACK to tell the user which B channel it is.
- The network sends the originating user a SETUP_ACK. This only acknowledges the SETUP. It carries no information about the outcome of the call.

It **<u>is</u>** mandatory to specify the Bearer Capabilities in a SETUP.

Although an initial SETUP would normally specify the destination number for the call, this is not mandatory. A SETUP is quite valid even if there is no called party number specified. This is how the Eicon DIVA LineCheck utility works. Where the destination phone number is available before the call is placed, for example, with Dial-Up networking, it is normal that the SETUP would contain this information.

In the case of a telephone, however, none of the digits of the destination number are known at the moment that the receiver is lifted. In this case the telephone sends a SETUP with any called party number information. The network may take this opportunity to attach a dial tone generator to the B channel. Alternatively, the telephone itself may generate a dial tone. This makes the ISDN telephone behave like a conventional analogue telephone.

As each number button on the telephone is pressed, an additional message called INFO is sent, which will report the number keyed. It is then the responsibility of the network to accumulate dialled digits until it has enough information to make a routing decision for the call. Naturally there has to be a limit to the length of time that the network will wait for the next digit. If this time limit is exceeded, the network will disconnect the (potential) call. This mode of operation is known as **overlapped sending**.

Where the complete called party number is known before sending the SETUP, the SETUP will also contain an indicator called **Sending Complete**. This tells the network that it should attempt the connection using only the digits supplied because no further digits will be sent by the TE. If the number is incomplete, this should be reported as an unknown destination error.

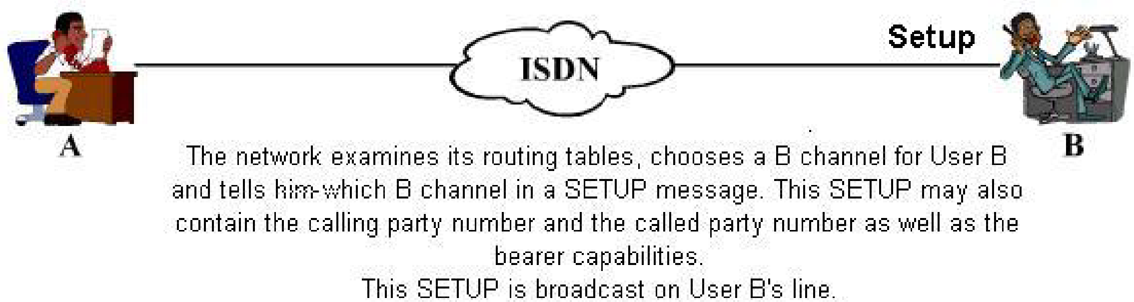

Q. Is the SETUP message sent to User B the same one that User A sent to the network?

Clue. Is it the same D channel?

Q. Is the SETUP message sent to User B the same one that User A sent to the network?

A. No, the SETUP message does not pass through the network. The network generates a new SETUP message to send to the User B. It may have copied some of the information that USER A provided.

Q. How many D channels are shown in the diagram?

A. The answer is 2. The D channel is not end-to-end. It only connects a user to the network.

More about the incoming SETUP

The incoming SETUP that indicates that a call is available will contain the call type (Bearer Capabilities). The called number (destination) and calling number (origin of call) are optional and may not be supplied.

If the called (destination) number is not supplied in the SETUP message then the implication is that the call can be answered by any device on the line.

Frequently, delivery of the calling (origin) number is an extra service from your service provider that you have to pay for. This service is often called CLIP (Caller Line Identification Presentation) or sometimes just CLI. See the section Managing Inbound Calls for more information on this subject.

At the destination, an ISDN device capable of receiving the incoming call responds with an ALERT message. An ALERT is sent to the originator of the call. This is just like the 'the remote phone is ringing' tone in a conventional telephone call.

In the case of an ISDN telephone, an ALERT means that the device has started to ring to alert the user to the incoming call.

When the device answers the call, it sends a CONN (connect) message to the network. A CONN is then sent to the originator of the call who then acknowledges it with CONN_ACK.

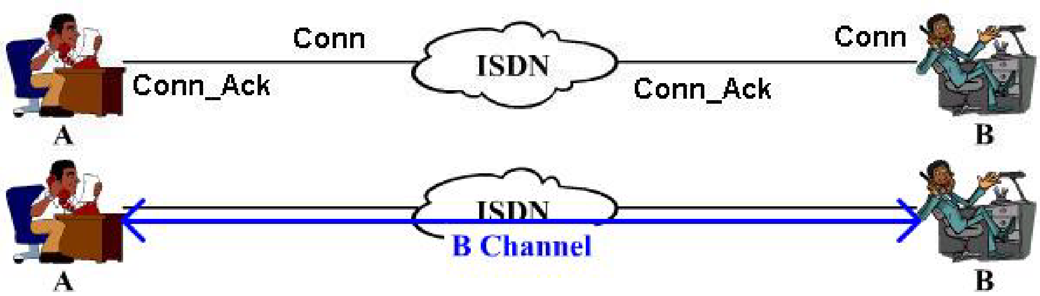

At the end of this process, a B channel has been established between the two users and data (which may, of course, represent voice or fax; or any other information) can flow freely between them.

The picture below summarises the exchange of messages during call set up.

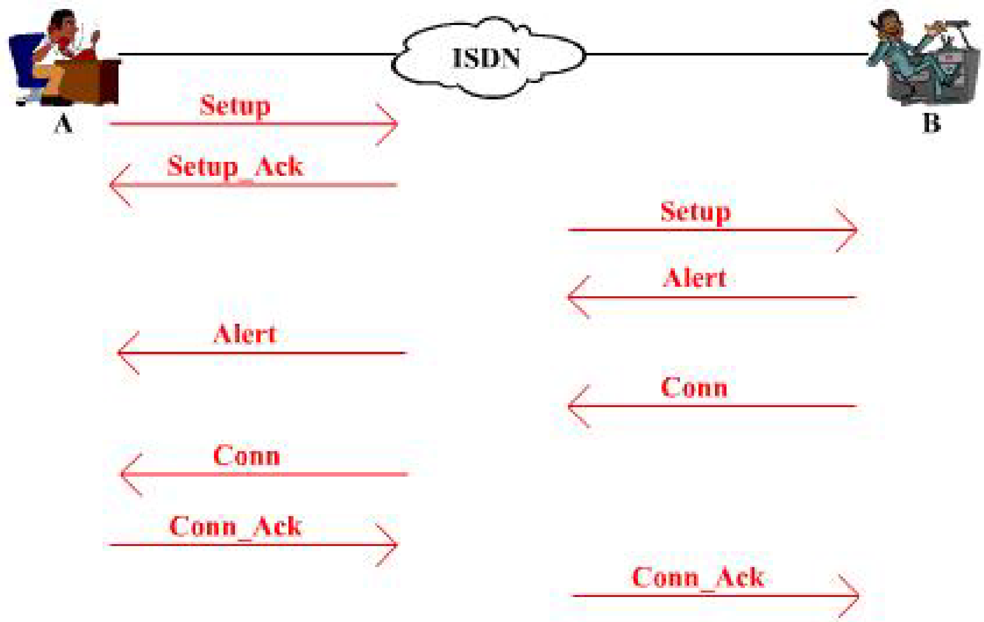

Finishing a call

Call disconnection also occurs through messages sent on the D-channel.

The message which conveys this information is the DISC, which also carries a cause code.

A DISC may occur because one of the users terminates the call (by hanging up the phone, for instance). The cause code will be 'normal call clearing'. A disconnection also may happen before the call has been completed, for instance with cause code 'destination busy', or for some other reason.

This diagram shows the interaction at the end of a successful call.

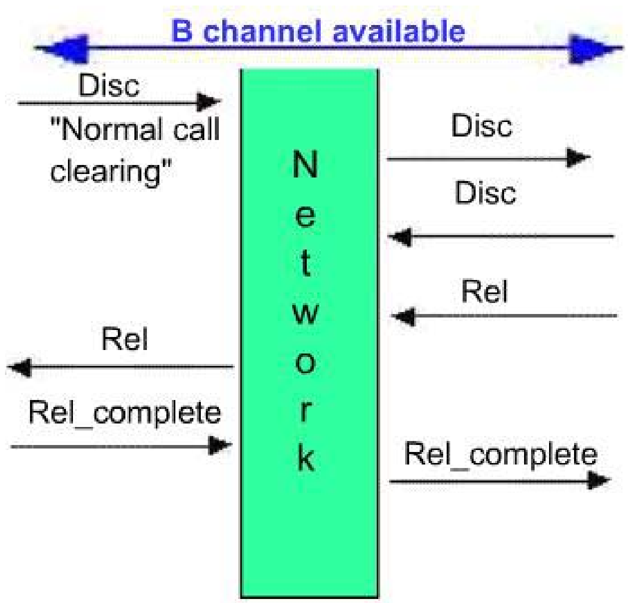

Receiving calls

When there is more than one device capable of receiving the incoming call at the destination, a mechanism exists to ensure that only one device can answer.

The SETUP is broadcast to all devices at the destination, and they in turn respond with an ALERT.

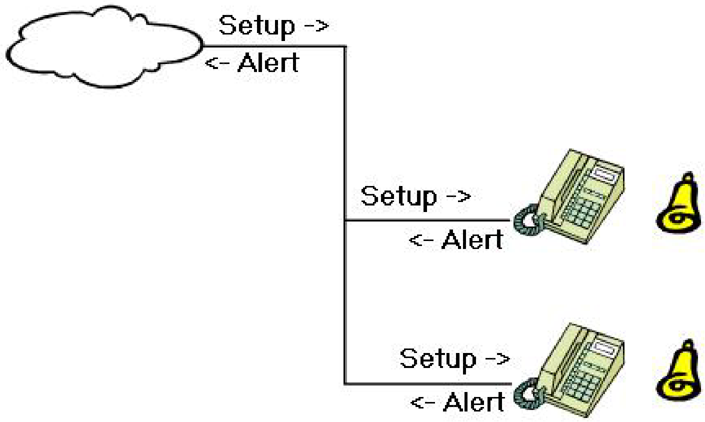

At this point, both devices are "ringing".
In this example, the upper device answers the call - a user picks up the telephone. This stops the bell ringing.

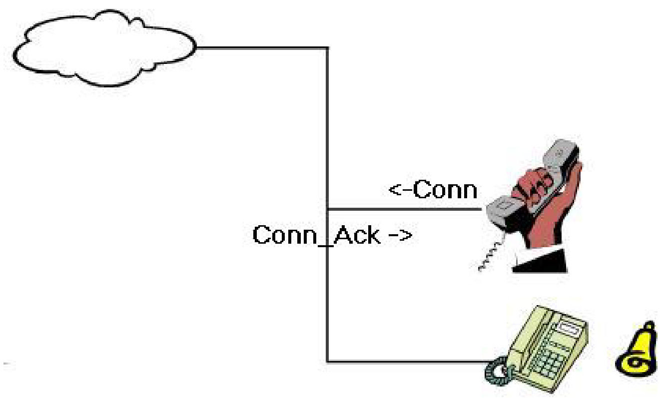

What about the lower telephone? Its bell is still ringing.

Stop other devices ringing

Since the network is aware of which device successfully answered the call, and also which devices have expressed an interest in the call, it is able to send a REL message to all those devices which sent an ALERT but did not answer. This stops these devices from ringing (or whatever is the equivalent activity for the type of device in question).

The interesting point about this scenario is that devices sharing the same line communicate with the network. They do **not** communicate with each other. In fact, they are unaware of the existence of each other.

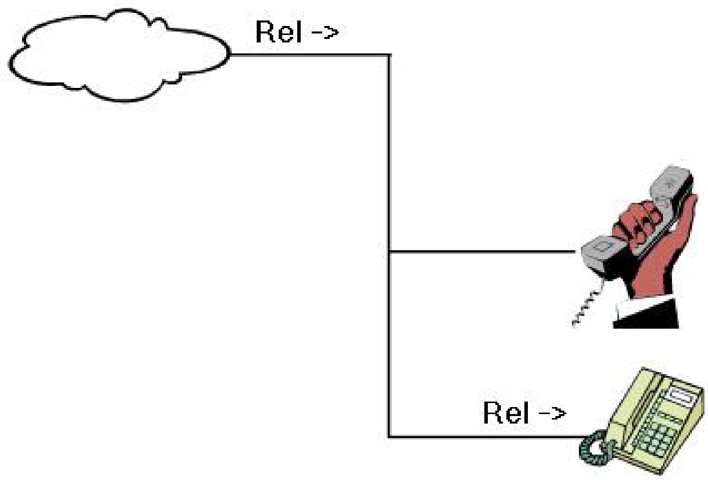

Connection collision

What happens if more than one device attempts to answer the call at the same moment?

The CONN messages collide in the D channel.
If you're wondering why we're not worried about the ALERT messages colliding (which could happen), this is because these messages are not critical to what happens next.

Don't forget that there is only one D channel and all devices that want to transmit have to contend for it. This is where the NT1 helps. It is able to signal to one or other device (or to both) that a collision has occurred.

It is not predictable which device will be informed of the collision, or whether they both are. The device which is not informed of the collision will not be aware that this happened and continues to transmit on the D channel successfully. The other device stops transmitting. It then waits for the D channel to become free and tries again. By this time, it's too late because the network has awarded the call to the device which won the collision.

Summary

In this section you learned:

- All calls are managed by a set of procedural messages between the network and the user
- These are all carried via the D channels at either end
- The end-to-end connection for the call itself is set up on the B Channel
- The most important messages are SETUP and DISC.
- Disconnections are accompanied by a cause code which will reveal why the call terminated
- Procedures also exist to manage multiple devices attempting to answer incoming calls
- Devices sharing a BRI do not cooperate with each other, they communicate with the network

Q/A

1. What channel do the call management messages use?
All call management messages are carried on the **D Channel.**

2. What channel does the end-to-end connection use?
End-to-end communications are carried on the **B Channel.**

3. What is the initial call request message called?
The initial call request message is a **SETUP.**

4. What is the call termination message called?
The call termination message is called a **DISC** (from DISConnect).

5. What information also accompanies this message?
The **reason for the call termination** accompanies the DISC message.

6. Why is this information significant?
Knowing the reason for a disconnect will allow ISDN communication problems to be diagnosed.

Managing inbound calls

In this section you will learn:

- How to manage more than one device on a single ISDN line
- How to ensure that the appropriate call types are directed to the correct type of device.
- What call filters are and how to set them

Handling incoming calls

Consider the example in the diagram. An ISDN line has a telephone, a fax machine, and a PC attached to it. The objective is to achieve the following:

- When someone rings the phone, only the phone rings.
- When someone rings the fax machine, only the fax machine answers.
- The PC never answers any calls.

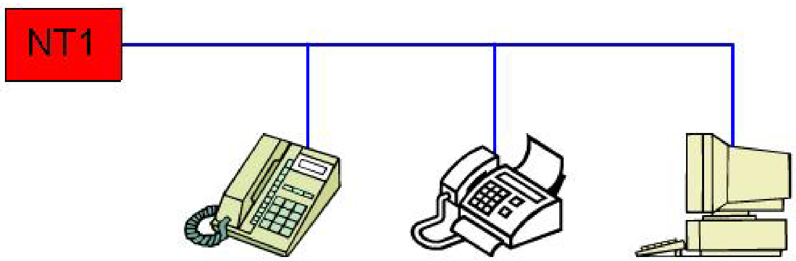

How might this be achieved?

When a device receives a SETUP message, it can apply a filter to that message. This filter examines the information that arrives in the SETUP message so that it may process or ignore the message. Only messages which pass the filter are acted upon.

In practice, equipment is only able to apply one test to each parameter. The kind of information which can be filtered in a SETUP message includes:

- Called party number (destination number)
- Calling party number (origin number)
- Call type (Bearer Capabilities)

ISDN calls can be divided into two main categories: analogue data and unstructured data. A SETUP **must** include the call type (bearer capabilities) of the call. Without this information, the SETUP is not valid and the network will instantly disconnect any such call request, citing "mandatory information message missing" as the cause. When a SETUP is sent by the network to the destination, the call type is also mandatory.

Handling incoming analogue calls

Within the analogue data call type, you can also distinguish between Group 3 fax and Voice.

So, the following filters could be set:

- Telephone only answers Voice calls, as indicated by the Bearer Capabilities
- Fax only answers Fax Group 3 calls, as indicated by the Bearer Capabilities

- PC doesn't answer because the are no programs running that are listening for incoming calls

However, there is a potential pitfall in this solution. This is because of the restriction that only one filter can be applied to each setup message.

A fax device, for example, may be configured to accept calls whose bearer capability is specifically Fax Group 3. However, fax messages which originate within an analogue network do not come with this very precise call type information. They are simply classified by ISDN as "generic analogue", just like a voice call.

Therefore, the fax device, which is looking specifically for fax calls, will reject valid fax calls from fax machines connected to the PSTN.

Similarly, if the fax machine is configured to accept "generic analogue" calls, then a fax message which originates directly from an ISDN network, and therefore may have the specific "fax call" type, will be ignored.

Consider the following case:

The fax machine on the analogue network makes a call to the fax machine on the ISDN network.

The SETUP generated on receipt of the call by the ISDN network specifies the call type is Generic Analogue. There must be a call type in a SETUP message, but the ISDN cannot know **exactly** what kind of device is causing this SETUP, so it makes a best guess.

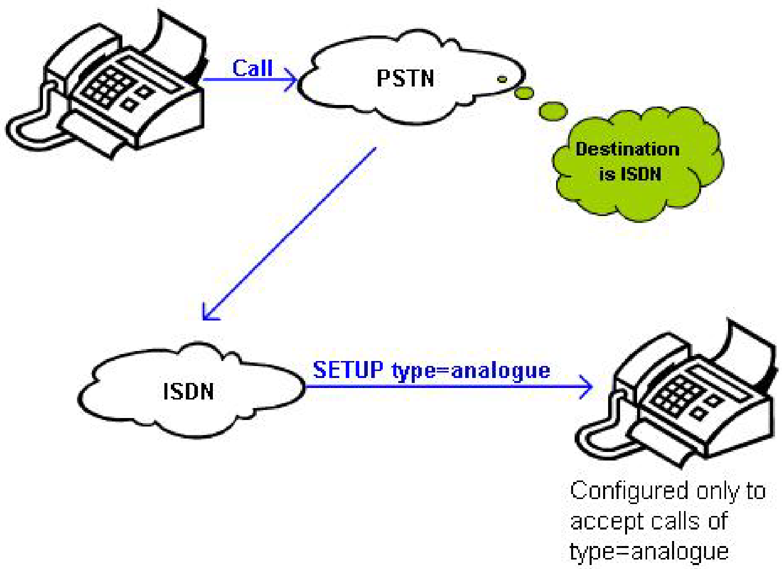

In this case, the fax machine in the bottom-right corner will happily answer the incoming call.

If however, a call is generated by a fax machine on an ISDN line, the SETUP message will describe the call type as being specifically from a "G3 Fax".

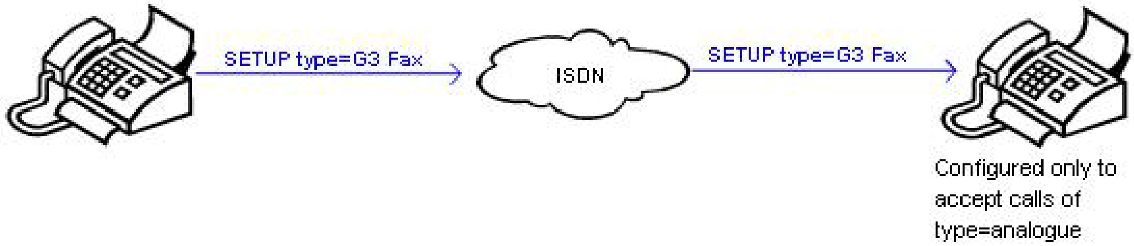

In this case, the destination fax machine will **not** answer the call as it is configured to accept only analogue calls.

In practice, a device can only apply one filter to each parameter of the incoming SETUP. A filter such as "CallType=G3Fax OR CallType=Analogue" is not normally possible.

For these reasons, this solution cannot be recommended. It would work very reliably if all networks were ISDN; and all Fax machines placed G3 Fax type calls and they were configured to answer such calls.

Fortunately, there is alternative method for directing incoming calls to the correct device - Multiple Subscriber Numbering, or MSN.

Using MSN for incoming calls

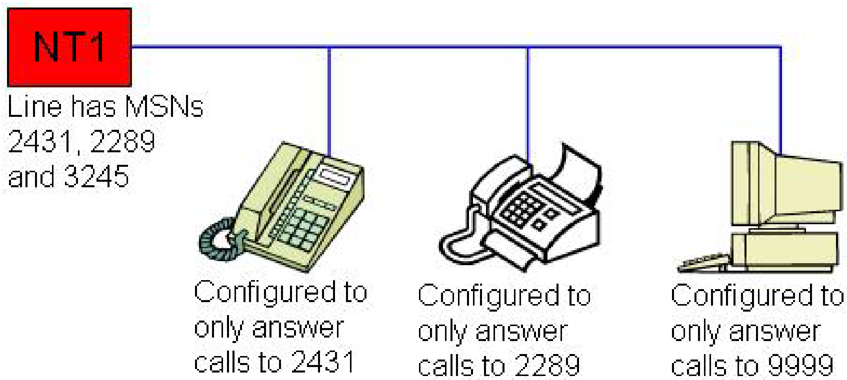

In this scenario we need a line with at least two subscriber numbers. This example has three numbers; the third number may be used by the PC to identify itself for outgoing calls.

The number belongs to the line, but is used by the device for filtering purposes.

This example also demonstrates a technique to stop a device from answering incoming calls. In this example, the PC is not running any programs that will answer incoming calls.

Let us assume now that the PC **is** running a program that answers incoming calls. We have no means of stopping the program, but we don't want it to answer any incoming calls. If we configure the PC to answer calls to 9999, the PC will never answer a call since this number will never be offered.

In practise, if you don't configure a Called Party Number filter, a device will not apply any filter on this information and answer the call provided it passes any other configured filters.

Summary

In this section you learned:

- SETUP messages can have a filter applied to them
- The Call Type can be used to filter calls
- The restriction to only one filter makes using Call Types impractical
- The most useful call filter is Called Party Number
- To make use of Called Party Number requires that a single ISDN line be assigned several different numbers
- This system is known as Multiple Subscriber Numbering or MSN

Q/A

1. What must each SETUP message contain?

Every SETUP message must also carry a **Call Type.**

This provides information about the contents - such as 'Group 3 Fax' or 'Generic Analogue'.

2. How many filters can be applied to a SETUP message?

Only **ONE** filter can be applied per SETUP message.

3. What call type is given to a fax call coming from an analogue network?

All calls originating on the analogue PSTN network will be automatically allocated the call type **'Generic Analogue'.**

A call might be fax or voice but it will always receive this general call type allocation because there is no equivalent ability within the PSTN to carry call type information.

4. A fax machine is configured to filter for analogue calls - will it accept G3 Fax calls?

No, it will only accept those calls whose call type is specifically 'Generic Analogue'.

This shows the shortcomings of using the call type filter to allocate calls to devices. Because there can only ever be one filter type, the results can never cater for all the possible combinations.

5. What filters are used when implementing MSN?

Called Party Number is the filter applied when implementing MSN.

The system relies on having M**ultiple Subscriber Numbers** per line and then allocating those numbers to a specific device as a filter. That number is then circulated as the number to call for that device, and it will answer only those calls which have that particular **called party number** as part of its SETUP.

6. Number of devices equals number of subscriber numbers required for MSN - true or false?

Not necessarily. If there are devices, such as a PC for example, which will never answer a call, then it is not necessary to have a subscriber number for that device.

The number of subscriber numbers required equals the number of devices which will be required to selectively answer incoming calls.

7. How might a PC be configured to ensure it never answers the phone?

A PC can be allocated a non-existent number. Since no call will ever arrive for that number, the PC will never answer any other calls.

Vairāk par MSM - http://www.eicon.com/support/helpweb/diva/msn.htm

Rate Adaption

In this section you will:

- Learn the difference between asynchronous and synchronous data streams
- Understand how they can be connected
- Examine the technologies available to achieve this
- Extend your understanding of the functions of an ISDN TA (Terminal Adapter)

Note: Rate adaptation is sometimes called 'rate adaption' in North America and some other countries.

Synchronous & Asynchronous data

The data stream on an analogue modem is intermittent because the traffic is **asynchronous** - that is to say, it starts and stops.

If you watch the lights on a modem when connected to an ISP, for example, you will notice that the transmit and receive data lights are off for much of the time.

On the other hand, B channel activity on an ISDN line never stops. This is because it is **synchronous**. This means that there is always data flowing.

There are two approaches that are used to handle the transition between these different types of data:

- **encapsulation** - the asynchronous data is carried across the digital network, and is then converted back to asynchronous data on the other side of the network.
- **conversion** - the asynchronous data is irretrievably converted to synchronous data. This is appropriate where the destination device can handle synchronous data; this method depends on an appropriate higher level protocol such as PPP. (Strictly speaking, this is not Rate Adaptation. However, Rate Adaptation is normally performed by terminal adapters, that are often also capable of performing protocol conversion.)

Encapsulation

The diagram below shows asynchronous traffic being transmitted across the ISDN network to a device that understands asynchronous data (e.g. an asynchronous port on a router providing Internet access at an ISP).

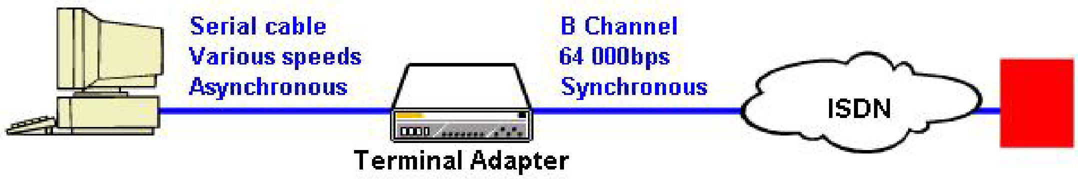

In order to allow the asynchronous traffic to pass across the ISDN, the terminal adapter (TA) between the PC's serial port and the network must fill all the empty gaps in the asynchronous data stream with something.

In addition, it must fill in the difference between the speed of the serial port and the speed of the B channel. This process is sometimes called **bit stuffing** because extra bits are inserted into the data stream to fill it out to 64kbps.

Bit stuffing

The bits required to fill up the 64kbps channel must be added in a systematic way so that the receiving device can remove the bits using exactly the same methodology to retain the information being transmitted.

There are two protocols commonly used that define systems to insert and remove extra bits:

- V.120
- V.110

You will use either one or the other, depending on the protocol required by the service you are calling.

Rate Adaptation is usually implemented in a Terminal Adapter. The activity happens within the B channel and is an end-to-end protocol. The ISDN is therefore unaware of this and the call type will be Unstructured Data.

Connecting 64k and 56k channels

The V.120 protocol is also used with unstructured data where you are connecting together B channels which have different information rates.

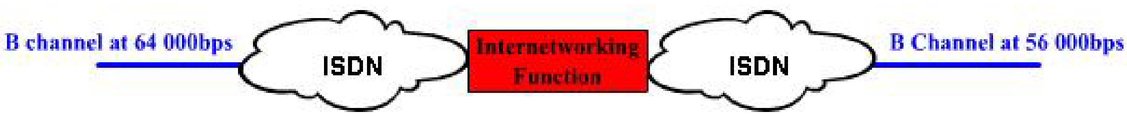

This is commonly used between Europe, where the B channel operates at 64 000bps, and North America, where some B channels operate at 56 000bps.

This is handled transparently by the ISDN providers at the network boundary points.

Connecting GSM data to ISDN

The V.110 protocol is often used when a mobile phone (GSM) connects to the ISDN in order to access a data service, such as a WAP gateway.

This conversion is handled transparently by the data gateway that connects the GSM network to the ISDN network.

Some GSM to ISDN gateways convert the call into a low speed modem format. Which format is used depends on the internetworking function chosen by the two network operators. The user has no control over this. Handling both types of call transparently is a challenge for the recipient of the call.

For voice calls between the GSM and ISDN networks, a different conversion is performed to convert the GSM format for voice calls to the ISDN format for voice calls.

Conversion

If you need to make an Internet connection using the serial port of your PC, but your ISP doesn't support V.110 or V.120, you will need to use **protocol conversion.**

Many Terminal Adapters support protocol conversion between **synchronous PPP** and **asynchronous PPP**.

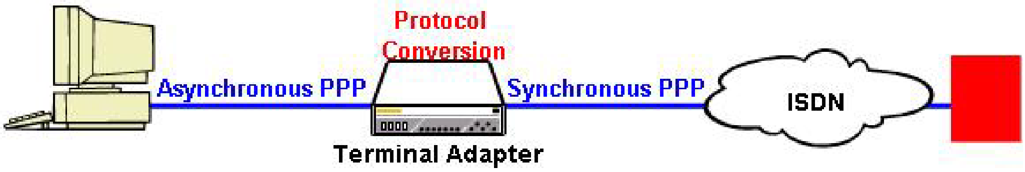

These protocols are not identical and conversion is required. The serial port of the PC is asynchronous so this determines the type of PPP used between PC and Terminal Adapter. The ISP, however, in natively supporting ISDN (rather than using Rate Adaptation), requires synchronous PPP.

Summary

In this section you learned:

- Asynchronous traffic contains gaps
- Synchronous traffic is continuous
- Asynchronous and synchronous data streams therefore need a conversion protocol between them
- V.110 or V.120 protocols are used to fill up the unused bandwidth
- This process is known as Rate Adaptation
- V.120 is also used to accommodate traffic flow between 56kbps and 64kbps ISDN lines
- Many Terminal Adapters also perform protocol conversion between synchronous PPP and asynchronous PPP

Q/A

1. Why is rate adaptation necessary?
Rate Adaptation is required when asynchronous and synchronous data streams need to be converted from one to another. This is because the synchronous steam is constant while the asynchronous stream contains gaps. Rate adaptation protocols fill the gaps and are also used to strip the gaps out again at the other end.

2. What are the two protocols which perform rate adaptation?
The two rate adaptation protocols are **V.110 and V.120.**

3. What determines which of these two protocols to use?
The selection of either V.110 or V.120 depends simply upon which of these protocols is being used at the other end of the connection.

4. Where else is rate adaptation employed?
Rate adaptation is also employed at the boundary points of different ISDN networks where the data rate is not the same.

For example, between the Euro-ISDN networks, which run at 64 kbps, and the US ISDN networks which run at 56 kbps, there will need to be V.120 rate adaptation.

5. What other method is there for converting synchronous and asynchronous data?
Another method of rate adaptation is called **Protocol Conversion.**

A typical example of this is when the Asynchronous PPP used by the serial port on a PC is converted to the Synchronous PPP used by ISPs.

6. Where does this conversion normally take place?
The protocol conversion takes place in a **Terminal Adapter.**

Cabling

In this section you will learn about

- Wiring requirements for BRI and PRI
- Maximum cable lengths
- How to connect several devices to a BRI
- When you need to use termination
- The power sources available through the interface

What you get with ISDN

EuroISDN is now the standard in most parts of the world, except North America and certain other countries.

The EuroISDN network is supplied to the user through a box mounted on the all. This box is called a Network Termination Type 1 or NT1.

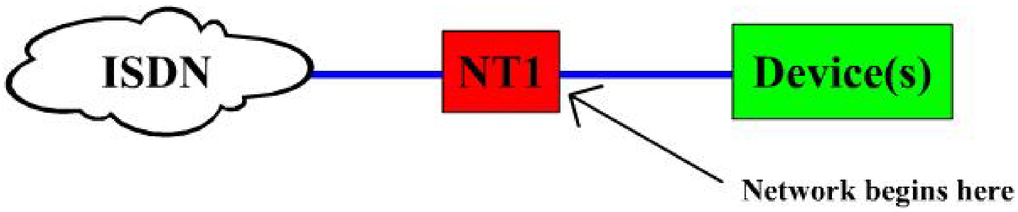

The NT1 is part of the network and is the point at which the network service is supplied.

In North America (and some other countries) ISDN is supplied as a passive socket terminating the circuit to the ISDN service provider.

Cabling from NT1

All the wiring on the user's side of the NT1 is the user's responsibility. Your ISDN service supplier is unlikely to be able to help you with any problems you have here.

In the case of the BRI NT1, the standard (I.430) is very specific. You must use RJ45 plugs (although the standard does not express it in exactly these terms). These are most commonly found in structured wiring systems based on Category 5 twisted pair. This standardisation is very important because you can connect as many as eight devices to a BRI NT1.

If your NT1 is providing a PRI, you will have to use whatever cables are suitable for the sockets that are provided. Unfortunately, these are can vary according to the ISDN supplier. They might be two miniature co-axial sockets, RJ45 sockets, or even screw-down connections. The standard that covers this (I.431) allows some leeway. This not usually a major problem, since you are allowed to connect only one device to a PRI NT1.

In the discussion below, all the distances cited assume that you are using Category 5 cabling. However, other types of cabling are also suitable such as AT&T PDS, Category 4, and Category 3.

In addition, these measurements apply only to the BRI. I.430 makes no mention of any distances at all - it only describes what the characteristics of the wiring should be. The lengths of the cables contribute to these characteristics. You may therefore see other values for distances quoted in other sources. Neither source is necessarily wrong, it just depends on the type of cable that you're using.

If you are connecting only one device to your BRI NT1, the cable can be as long as 1 km if you are using Category 5 twisted pair.

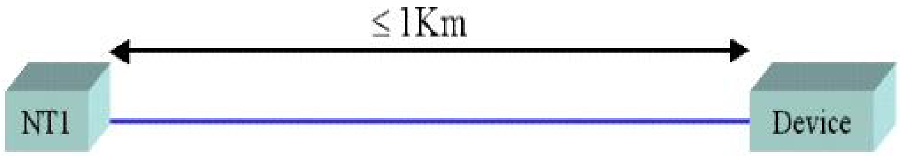

In practice, if you really need to run a cable for this distance it would probably be safer to re-site your NT1 closer to the equipment to be connected.

All NT1s have two sockets for devices, so you could add another device to the other socket at a similar distance.

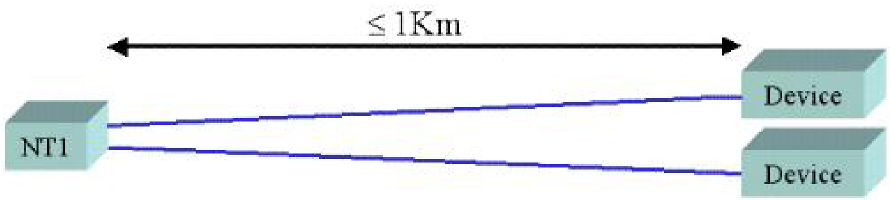

This form of wiring is known as **point-to-point**.

As many as eight devices can be connected to a BRI NT1, but the NT1 normally only has two sockets. You must use some special wiring to connect more than two devices and there are also rules about the maximum distances involved.

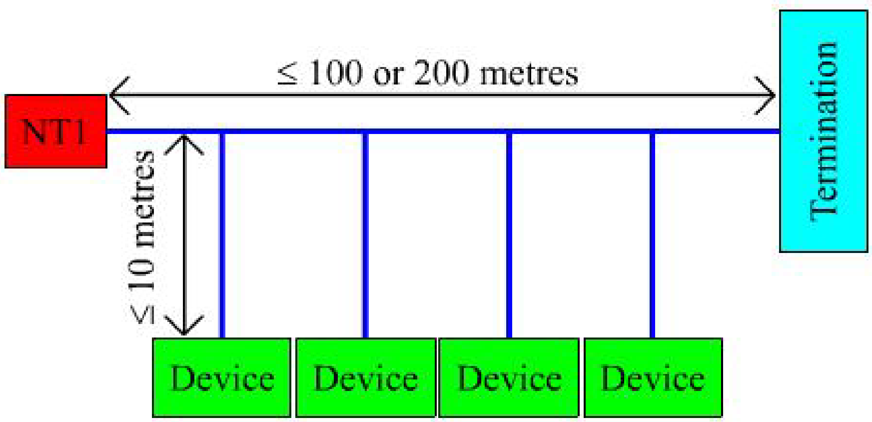

In this configuration, known as the **short passive bus,** the devices are randomly distributed along the length of the cable. Each of the drop cables (between the main cable and the device) may be as long as ten metres.

The total length of the main cable may not be longer than either 100 metres or 200 metres, depending on the quality of the cable. Cable has a characteristic impedance, either 100Ω or 200Ω. With the 100Ω cable, the maximum length is 100 metres, whereas with the more expensive 200Ω cable, the maximum length can be 200 metres. The characteristic impedance is usually quoted by the cable supplier.

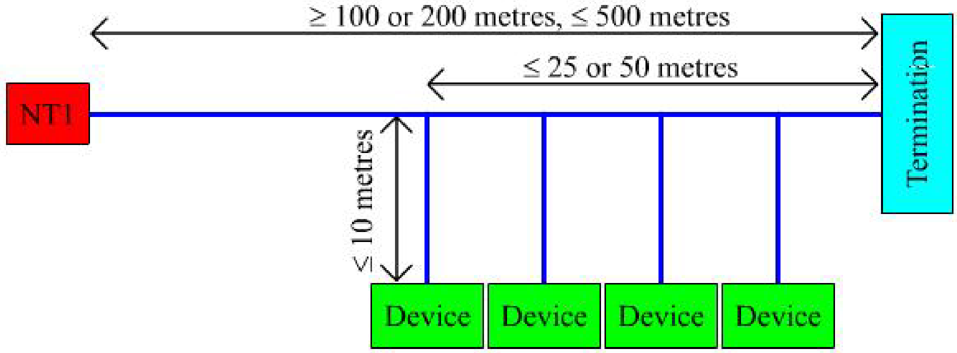

Another form of wiring is known as the **extended passive bus**. In this case the total maximum distance allowed is greater, up to about 500 metres, but there is a restriction on where the devices may be located on the main cable.

The maximum length of the drop cable is the same as before (ten metres). The restriction is that the device nearest the NT1 must be no more than either 25 or 50 metres from the end of the main cable furthest from the NT1. The smaller distance is for 100Ω cable, the larger for 200Ω cable.

The cable for the BRI can contain as many as eight conductors. However, four of these are optional. They are organised in pairs like this.

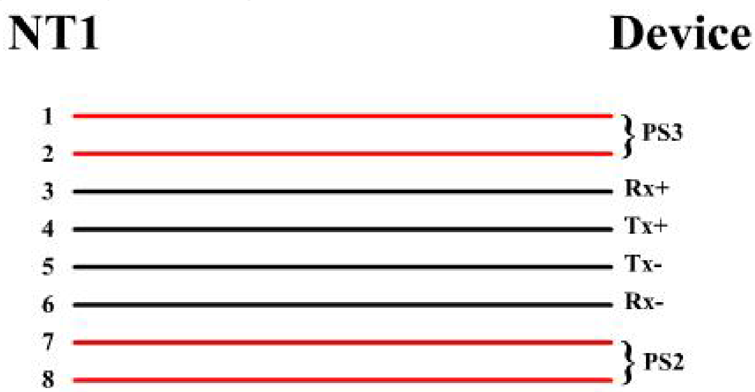

Notice how there is one pair of conductors which carry data away from the devices toward the NT1 and another pair which carry data from the NT1 toward the devices.

The red lines show the optional pairs. These are concerned with providing power to devices connected to the BRI. This is a convenience for an ISDN telephone, that, for example, has an LCD and LEDs. This therefore needs an external power supply which can be supplied through its ISDN interface rather than through a separate mains power adapter.

Power provided by some NT1s

PS2 is a power supply provided by the NT1. Not all NT1s can do this as standard. For instance, in the UK this is a separate, chargeable feature.

PS3 is a mechanism through which one device can supply power to other devices. This is very rarely implemented.

There is another power supply known as PS1. This is also known as the "phantom" power supply. It is found in the voltage difference between the transmit and receive data pairs. This derives from the network itself and is normally used for emergency situations where the user's local power supply has failed. Contrast this with PS2 where the NT1 provides this power from its own mains plug.

If none of the equipment that you have connected to your NT1 requires a power supply through the interface, you do not have to plug the NT1 into the mains.

The NT1, which is not a passive device, normally powers itself from PS1. Only if the NT1 contains a TA (so it is a 'NT1-plus' or 'Super-NT1') or if it can supply power to attached devices will it need to have a mains power supply.

Some ISDN phones can also be powered from PS1. This allows you to have at least one working ISDN telephone on your line in the event that your local power supply fails. In some countries this is a regulatory requirement.

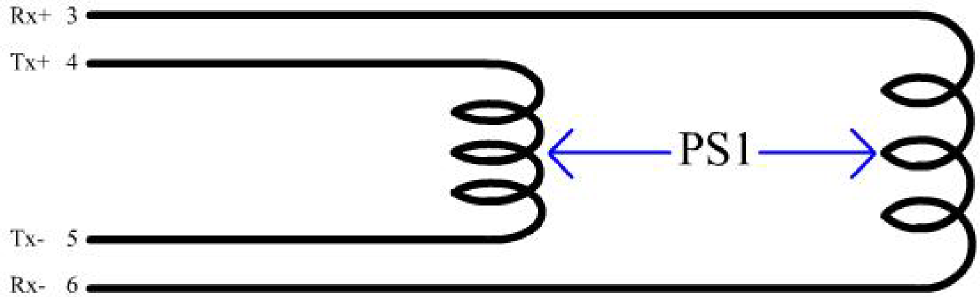

Note that there are no provisions for PS1, PS2 or PS3 in the case of the PRI.

Termination

If any of the cables attached to an NT1 is longer than 75 metres, then **both** ends of of **both** cables require termination.

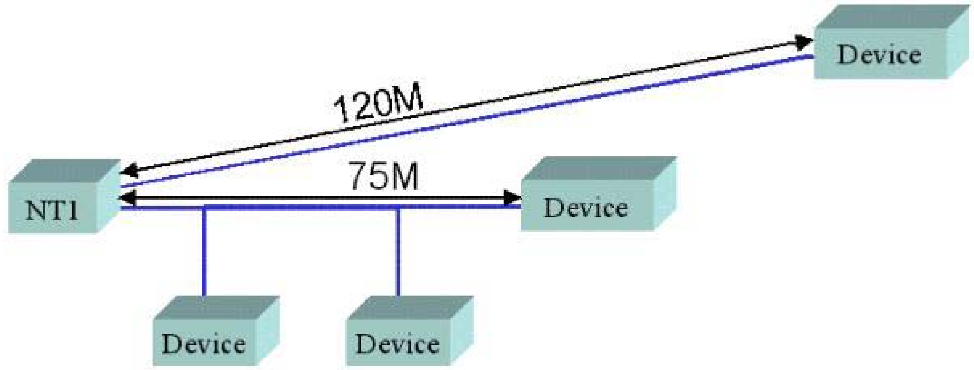

Terminating the NT1 end of the cable is simple. All NT1s have at least one switch that is concerned with this. It will be marked "short" and "long" (or the equivalent in the local language).

Termination at the far end of the cable involves installing two 100Ω resistors across each of the pairs carrying signals. It is done like this.

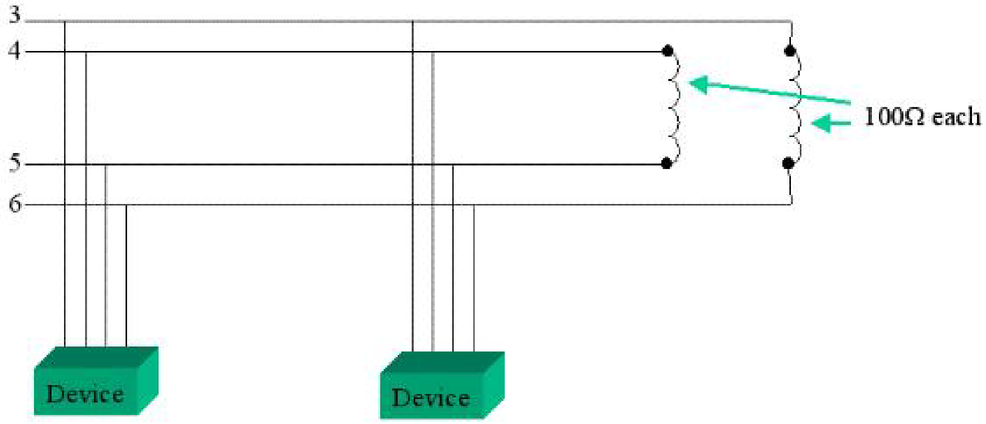

This also shows how you arrange the wiring for the short and long passive bus configuration. If you also need the power supplies PS2 and PS3, you should include conductors 1, 2, 7 and 8 as well.

Summary

In this section you have learned that:

- The NT1 is an integral part of the network in EuroISDN.
- The user provides the NT1 in North America.
- You can use most types of twisted pair cabling for a BRI.
- The type of cabling used for a PRI is not consistent and depends upon the local provider
- For a BRI you can use point-to-point wiring, or point-to-multipoint wiring, as either a short passive bus or an extended passive bus.
- A BRI may be able to provide a power supply through the interface.
- Any cable over a certain length needs to be terminated at both ends.

Q/A

1. What is the purpose of the switch that all NT1s have?

This switch controls **termination** and should be activated if any of the cables being used to attach devices to the ISDN is greater than 75m long.

2. If none of the equipment connected to an NT1 requires a power supply through the interface, do I need to attach the NT1 to the mains?

No - there is sufficient power within the ISDN standard to power the NT1 itself.

3. How many devices may be connected to a BRI?

Up to **eight** devices may be connected to a standard Basic Rate interface.

4. How many devices may be connected to a PRI?

Only **one** device may be connected to a Primary Rate interface.

5. What kind of cabling is used for a BRI?

Normally, it will be the same as Category 5 cabling with an RJ45 jack.

6. What kind of cabling is used for a PRI?

The cabling standards for the Primary Rate interface are not as specific as they are for Basic Rate, and they will therefore vary according to the local ISDN supplier.

More Termins

http://www.freesoft.org/CIE/Topics/index.htm

Analogue

What does "Analogue" mean?

Often this refers to a quantity that is can vary continuously. That doesn't mean that it's always changing. It means that it can have any value.

One example of this is temperature. The temperature outside today might be 25°C or it might be 35°C. It could also be any other temperature, perhaps 33.5°C or 36.1274°C*.

If you were a farmer, you might have six cows or seven cows, or 200. You wouldn't have 11.5 cows unless you're called Damian Hurst (an artist noted for exhibiting animals that have been split down the middle).

The number of cows you have can only be discrete values, while the temperature can be any value.

Baseband and Broadband

Baseband communication is fundamentally digital. Data is transmitted directly as on/off signals in a transmission medium (usually a copper wire). Because the transition from on to off and vice versa is usually very nearly instantaneous, a large bandwidth is required. Generally, only one communication channel is available at any given moment.

An alternative to this is broadband communication which (paradoxically) requires a lower bandwidth. In broadband communications, data is carried by a particular portion of the available bandwidth. This allows more than one communication channel to be carried a time. This is comparable to the way radio stations can transmit on different frequencies so that they don't interfere with each other.

PBX or PABX

PBX stands for Private Branch Exchange. PABX means Private Automatic Branch Exchange. Nowadays, these both mean the same thing.

This is a telephone exchange as would be used within an office to connect a number of internal users to a (usually) smaller number of outside lines on the public network. They can range in size from very small (two outside lines and maybe eight extensions) to very large (hundreds of outside lines and thousands of extensions).

This expression covers both conventional (analogue) exchanges and those supporting ISDN.

Transmitting data over the D channel

There are two ways in which the D channel can be used to carry data between two users:

Some networks can optionally connect each user's D channel to an X.25 data network. Then, data can be transmitted on the D channel as far as the ISDN network, where the data is re-directed onto the X.25 network.

The ISDN specifications allow for a small amount of user-data to be attached to the SETUP message that initiates the call. This could be used for applications requiring very short exchanges of data (such as credit card authorisation).

Protocol

A protocol describes a procedure.

When you declare your taxes, you complete a form. There will be boxes on this form where you state your name, your income, deductions, etc. This makes it easier for the person (or machine) that has to process the form; they will know exactly where to find each item of data.

There is also a set of rules which will describe what happens to the form inside the tax office. If you declare income from another country, it's likely that your form will be sent to different department from those forms where no such income has been declared.

Data communications are very similar. You use a protocol to communicate information to another person (or computer). If you both follow different rules about how the information will be laid out, you won't be able to understand each other. The communications protocol will also dictate the actions that must be performed given the information received.

SNA

SNA stands for Systems Network Architecture.

This is a family of protocols defined by IBM over the last 30 years for use with their mainframes and mid-range computers. There are also third party implementations (including Eicon Technology) of some of the protocols on PCs. It is widely used in large organisations such as banks, public utilities and public administration for creating large networks.

PPP

PPP stands for Point to Point Protocol.

This protocol is widely used for dial-up access to the Internet. All communications on the Internet are carried by a family of protocols called the internet protocol suite. One of the features of PPP provides a mechanism for reliably transporting the internet protocol suite across unreliable telephone lines.

PSTN or POTS

PSTN stands for Public Switched Telephone Network. POTS stands for Plain Old Telephone Service.

Both of these terms refer to the conventional analogue telephone system. Often these expressions are used to make contrast with ISDN.

MLPPP

This stands for Multi-Link Point-to-Point Protocol.

This form of PPP is able to coordinate data across several parallel connections. It ensures that all items of data have been properly assembled in the correct sequence at the receiving station before passing them onward.

It is frequently used across two ISDN B channels, but can be used with any form of data communications, such as modems; and across any number of data links.

When using MLPPP, it is advisable to use data links which are of approximately the same speed (eg. 2 B channels; or 1 B channel and a 56K modem). If the data links have greatly different speeds, the data throughput will be limited by the speed of the slowest link. This is because the receiving station will always be waiting for data from the slowest link to complete some of the logical units of data, before they can be passed onward.

Remember that MLPPP is not the same two separate PPP links. MLPPP is a single connection which uses more than one data link. Two PPP connections are two separate connections to two independent destinations.

PBX or PABX

PBX stands for Private Branch Exchange. PABX means Private Automatic Branch Exchange. Nowadays, these both mean the same thing.

This is a telephone exchange as would be used within an office to connect a number of internal users to a (usually) smaller number of outside lines on the public network. They can range in size from very small (two outside lines and maybe eight extensions) to very large (hundreds of outside lines and thousands of extensions).

This expression covers both conventional (analogue) exchanges and those supporting ISDN.

ISPBX

This stands for Integrated Services Private Branch Exchange.

This is a telephone exchange as would be used within an office to connect a number of internal users to a (usually) smaller number of outside lines on the public network. They can range in size from very small (two outside lines and maybe eight extensions) to very large (hundreds of outside lines and thousands of extensions).

The public network lines for an ISPBX must be ISDN. The extensions may be served by an internal ISDN or a conventional analogue telephone network, or both.

ISPBXs are a source of many ISDN problems, because they may be configured in so many different ways and because the ISDN implementation is sometimes incorrect. When diagnosing ISDN connectivity problems, it is always best to test your connection on a public ISDN line before calling for support.

Group 4 Fax

Group 4 fax is a method for sending documents at high speed. It only works on ISDN lines and can only communicate with another Group 4 Fax machine. In contrast with Group 3 which might need as few as 30 seconds to transmit a single page, Group 4 needs only 3 to 5 seconds for the same page.

Group 4 fax machines are quite expensive and very few are installed compared with the number of Group 3 (analogue) fax machines. There are also a number of software packages which emulate a Group 4 fax machine.

Bearer Capabilities

This is a mechanism by which a user can inform the ISDN of the protocol that he is using in a B channel.

Strictly speaking, the Bearer Capability is one of three mechanisms that are used to communicate this information. However, the other two mechanisms (called High Layer Compatibility and Low Layer Compatibility) are always used in conjunction with Bearer Capabilities, so they are often known collectively as bearer capabilities. The High Layer Compatibility and Low Layer Compatibility are used to provide additional details about the B channel protocol. On this page, Bearer Capabilities (capitalised) refers to the mechanism, bearer capabilities (lower case) refer to the protocol in use.

When a user requests a connection through an ISDN, he is obliged to specify the Bearer Capabilities. If he does not, his request is rejected. If he is using a protocol that the ISDN does not understand he will specify that the data is transparent (or unstructured). High Layer Compatibility and Low Layer Compatibility are both optional.

The bearer capabilities are communicated to the destination, so that he may act upon them. For example, an analogue call could be presented to a modem, but a transparent data call would be presented to an ISDN adapter.

The ISDN might also use the bearer capabilities for its own purpose. For example, in North America, some network operators make different charges for voice and for data calls. Some networks are able to optimise their internal performance by using the bearer capabilities.

PSTN or POTS

PSTN stands for Public Switched Telephone Network. POTS stands for Plain Old Telephone Service.

Both of these terms refer to the conventional analogue telephone system. Often these expressions are used to make contrast with ISDN.

DSP (Digital Signal Processor)

A digital signal processor a high-performance microprocessor that is optimised to provide real-time of analogue data. They typically have a small instruction set, but can perform many millions of floating point operations per second.

For example, the are capable of performing the Fast Fourier Transforms necessary for sampling analogue signals input and creating a bit stream representing the bit stream that represents this, and vice versa, in real time.

Start-Stop Protocol

In the 1960s and 1970s, IBM used to call asynchronous traffic by the name Start-Stop Protocol. This is because there are not always bits present in the communications link.

In such a protocol, the data is divided into bytes. These bytes may be of any length, normally between four and eight bits, but all will be the same size.

Each byte is followed by one, one-and-a-half or two stop bits. These are used as an inter-byte filler and allow the receiving modem to re-synchronise its bit timing. This is especially important when one byte immediately follows another. Following a gap in transmission, the receiving modem is able to re-establish bit synchronisation at the moment it receives it receives the first bit of next byte.

This means that an asynchronous channel has three states:
Transmitting a zero
Transmitting a one
Not transmitting

By contrast a synchronous channel has only two states:
Transmitting a zero
Transmitting a one

Bit Rate and Information Rate

There is a distinction between these two ideas. The bit rate is the advertised bit rate of a communications channel. In the case of an ISDN B channel in Europe it is 64 000 bits per second.

Information rate is the rate at which information is passed over the channel. The information rate can never be higher than the bit rate, but it might be lower.

An ISDN telephone transmits and receives 64 000 bits per second, even when there is total silence.

Synchronous and Asynchronous PPP

Synchronous and asynchronous PPP are basically the same, but they are not compatible with each other. Both can be used to send binary data.

One of the chief difference is that asynchronous PPP can be sent across an asynchronous interface. Such interfaces have a flow control mechanism used to regulate the flow of data so that no component (eg. buffer, modem) is overwhelmed. There is often a choice of flow control: hardware or software. This is demonstrated in the panel below provided by Microsoft on Windows 98.

The software flow control looks for two characters in the data stream. These are the control (non-printable) characters XON and XOFF. XOFF stops the interface from transmitting until it sees an XON character. The data that asynchronous PPP is transporting might contain these characters (PPP is bit-transparent). This could have disastrous consequences on the transport of the data.

Asynchronous PPP uses a technique called byte stuffing. If a control character is found in the data stream that might have some significance to the data link, an extra character (call an escape character) is inserted before the character in question. The control character is converted into something else using a table which maps control characters into something innocuous to the data link.

When the data is received, the escape character is removed and the reverse mapping is performed on the following character. This recovers the original data. There is an overhead to this process which, depending on the data being transmitted, could be as high as 50% additional data.