Speaking Clearly with ATM:

A practical guide to carrying voice over ATM

Voice and ATM - What are the Issues & Objectives

Why run Voice over ATM at all?

Given the commercial pressures facing any network operator, there has been a continuing drive to reduce operating costs and lift network efficiency. ATM is one of the principal tools in the network designer's armoury to achieve these goals.

ATM allows the statistical multiplexing of traffic over any network resource. Statistical multiplexing avoids the need to pre-allocate resources for a user, but allocates resources on-demand, providing bandwidth to users when, and only when, they need it. This allows the network to support more users (typically twice as many as a simple TDM network).

However, to succeed an ATM network must be able to integrate traffic of all types into this single integrated, statistically multiplexed stream. This means handling that most testing of traffic - Voice. The transport of voice presents the network with a number of technical challenges; these must be successfully managed if the network is to provide an effective voice transport mechanism.

Today, through the work of the ATM Forum and its members, these issues have been addressed and it is possible to build and operate an ATM network to meet the needs of many types of voice application.

This paper looks at the requirement to carry voice over an ATM network, the technical challenges to building such a network, and the solutions that have been developed.

Today, it is possible to build an ATM network to meet the VPN voice needs of a corporation, or the backbone transmission needs of national voice network, or to support a national cellular network, or provide international voice transport. This paper also analyses the differing requirement presented by each type of application.

However, we must start by analysing the basic requirements placed on any network that is to carry voice traffic.

What’s the Problem?

Why is voice so tough?

Perhaps because voice and video are so closely linked to our everyday senses they provide a severe trial of any networking technology, and ATM is no different in this respect.

Any packetised approach to the carriage of voice faces a number of technical challenges. These challenges are the same for any packetised technology, be it ATM, Frame Relay, IP or X.25.

These challenges spring from the real-time and interactive/transactional nature of voice traffic. To understand how these problems have been addressed by current ATM technology, we first must understand what they are, and how they arise.

Delay

The first issue is delay. In order to hold an interactive conversation end-to-end, network-induced, delay must be kept to a minimum to ensure that subscribers receive an acceptable quality of service.

There are two significant thresholds that are crossed as network delay increases.

The first of these thresholds results from an electrical quirk of the world's telephone network. Telephone handsets use 4 wires to connect to the network. However, the long distance network only uses a single pair of lines to carry a telephone call. This transition from 2-wire to 4-wire operation is accomplished by the use of a device known as a "hybrid" at each end of the network.

Hybrids however cause a reflection of signal, due to unavoidable impedance mismatch. The effect of this is that some of your speech signal is reflected back towards you by each hybrid, and arrives in the earpiece of your handset. Normally you do not notice this, as it just like hearing your own voice through the air of the room. However, if there is delay across the network, the reflected signal from the far-side hybrid will arrive back in your earpiece as an echo. The effect of this echo is more pronounced as the delay increases.

At about 30 milliseconds of network delay, the echo is so significant as to make normal conversation very difficult. Therefore, once a voice circuit delay exceeds 30 milliseconds we must include echo cancellors within the network. These devices are expensive and complex. Moreover, they operate most successfully when the delay over a circuit is constant, and as we shall see this may not be the case in a packetised voice network.

Once in place echo cancellation systems allow network delays to reach approximately 150 milliseconds before further voice quality degradation is experienced.

At about 150 milliseconds of delay a second problem emerges. At this level of delay, we start to have significant problems in carrying on a normal conversation. Normal conversation patterns demand that some responses from the listener are received within less than 200 milliseconds, and delays of this order result in stilted conversations, and "clashing" (where both parties try to talk at once. This type of problem is commonly encountered over satellite links.

For these reasons toll quality networks, require that the end to end network delay for voice traffic is less than 25 milliseconds in national networks and less than 100 milliseconds in an international context.

Where does the Delay come from?

It is not immediately obvious why ATM networking results in delay problems — after all ATM is very fast… The difficulty stems from three areas:

- packetisation delay,
- buffering delay,
- encoding delay.

Moreover, these three issues are common to all packetised networks, irrespective of transmission speed.

Packetisation Delay (aka Cell Construction Delay)

Packetisation delay is the delay caused by the necessity to fill a packet or cell before it is transmitted.

Voice samples are received at a rate corresponding to the compression level of the voice traffic. Therefore, normal PCM encoded voice arrives at a rate of 64Kbps. Compressed voice will arrive at lower rates (32Kbps, 16Kbps, and 8Kbps). A delay is incurred while sufficient encoded data is received to allow the packet or cell to be filled for transmission.

This delay is clearly defined by the minimum packet length that can be dispatched, and the encoding rate. While in some packetised schemes it is possible to send very small packets, this is unattractive as the efficiency of any packetised scheme falls as the ratio between payload and header length falls. In the case of ATM its fixed cell length means that delay is just proportional to the level of voice compression employed — greater compression resulting in greater delay.

To take a practical example, if we are using 8kbps LD-CELP compression and filling ATM AAL1 cells, one cell is filled every 47 milliseconds. As the cell cannot be transmitted until it is full, we have broken our first delay threshold before the traffic has even been inserted into an ATM cell!

This problem can be addressed either with partially filled cells (somewhat inefficient) or by multiplexing several voice calls into a single ATM VCC.

The second area of delay results from the need to maintain a real-time delivery of voice traffic across the network. When traffic is presented to an ATM (or other packetised network) it must first be broken into small units (packets, frames or cells) for transmission. At the destination these must be reassembled into a facsimile of the original voice call (the emulated circuit). In the case of voice, the traffic must be reassembled in real-time with no temporal distortion.

Buffering (Build-Out) Delay

To operate correctly the re-assembly function of any constant bit rate (CBR) or real-time data stream clearly requires that the data required to reconstruct the carried traffic is available at the destination at the correct moment in time. If a cell is delayed in transit, then the reassembling (SAR) function may "under-run" (having no data to process). This would result in gaps in conversation.

In order to prevent this from occurring the receiving SAR function will accumulate a buffer of information prior to commencing the reconstruction of the traffic from received cells. The depth of this buffer must exceed the maximum predicted delay propagated by the network to ensure that no under-runs occur.

However, the depth of this buffer translates into delay, as each cell must progress through the buffer on arrival (at the emulated circuit's line rate).

The implication of this is that the Cell Delay Variation within the ATM network needs to be tightly controlled to allow the minimum depth of reassembly buffer to be configured, and thus minimise the network delay.

It important for this reason to see that Cell Delay Variation in a network can have a major impact on total network delay and is a major criterion in the design of ATM equipment.

Encoding Delay

The third source of delay occurs in the encoding of the analogue signal into a digital form.

The desire to compress voice traffic to the minimum bit rate while maintaining quality results in processing delay. This typically becomes more significant at lower bit rates. The family of Low Delay Code Excited Linear Prediction (LD-CELP) encoding algorithms are now quite widely used and allow toll-quality encoding to be achieved at bit rates as low as 16Kbps with encoding delays of less than 10 milliseconds. However, as we have seen, this may represent 50% of our total delay budget.

Signalling Support

The second set of issues to be addressed relate the efficient utilisation of resources and the transfer of control and signalling information.

A voice call consists of two parts - the voice samples and the signalling information. The signalling information includes the dialled number, the on-hook/off-hook status of the call, and possibly a variety of other routing and control information. This signalling can be encoded in a number of ways, and may be sent as common channel, channel associated, or DTMF dialled digits.

Typically, multiple voice channels are combined into a single circuit. So for example, a European 2Mbps E1 circuit contains 30 discrete voice channels, or an American 1.5Mbps T1 circuit contains 24. Signalling information may be embedded within each discrete voice channel (known as channel associated signalling) or aggregated into a single signalling channel, containing signalling information for all the channels on the circuit (common channel signalling).

In the simplest case, where voice traffic is being "trunked" across an ATM network, we may just wish to transfer the signalling end to end. However, in more complex examples using, common channel signalling, where traffic from one site must be switched and delivered to two or more endpoints, simple approaches may not work. In these cases the signalling channels must be terminated and interpreted at the ATM switch so that the correct information can be passed to the correct endpoint. Given the fact that many signalling mechanisms (both proprietary and standardised) exist, the ATM network may need to understand several types of signalling protocol.

A secondary opportunity exists for the intelligent use of signalling in an ATM environment. The on-hook/off-hook signals can be used to allocate active voice channels to active VCCs within the network. In this way, an ATM network can ensure that valuable bandwidth is only allocated to active calls, and a large community of users can contend for available resources, as only active circuits will be transmitted.

Synchronisation

Synchronisation of traffic must also be addressed. ATM, as its name implies, is asynchronous in nature, but the transport of voice demands that the data be synchronised to maintain the temporal relationship between speaker and listener. With simple point to point applications, the two end-points may be synchronised by two standardised mechanisms: Adaptive Clocking or

Synchronous Residual Time Stamping (SRTS). These adjust the clock rate at one end of the circuit based on the clock rate of the other end.

Adaptive Clocking does this by monitoring the depth of the SAR receive buffers. It then adjusts the clock rate of the "slave" end to maintain an appropriate buffer depth.

SRTS monitors the rate of the line clock at the "master" end of the circuit, with respect to a standard clock. The difference between the two clocks is then encoded and transmitted as part of the data stream. At the slave end this difference signal is retrieved and used to adjust the slave clock by reference to the difference signal and the standard clock.

However, neither of these mechanisms can operate where multipoint services are in operation. For example in the SRTS case it is clear that is not possible for a slave site "A" to adjust its clock in response to two difference signals being received from two different master sites. Neither, in the adaptive case, can a slave adjust its line clock in response to changes in the buffer depth resulting from two different data sources without adversely affecting one of them.

Therefore, for multipoint services we must adopt an externally synchronised model where each node in the network is synchronised to some external clock source. In practice, this is easy to achieve given the availability of global timing standards.

The Applications

Who needs ATM voice?

There are many applications for the transport of voice over an ATM network. These encompass both the enterprise network builder, and public service provider. Each application has differing requirements; this section considers the applications and the requirements they have for voice transport.

For this discussion three classes of network operator are clearly defined:

- National or International Operators.
- Alternate Carriers or Value Added Network Suppliers
- Enterprise (Private) Networks

Let us first discuss the characteristics of each of these groups.

National or International Operators.

These organisations will typically have an extensive PSTN service in place operating over existing PDH or SDH/SONET infrastructure. They will additionally have data networking services supporting business and residential users with multiple, discrete networks supporting these applications.

Where bandwidth is limited there will be a strong requirement to integrate voice and data traffic into a single ATM network for reasons of efficiency. This is particularly true within the international segment of a carrier's network where the cost of leased international bandwidth demands top levels of efficiency

Within the local loop, ATM may be a valuable solution for the carriage of voice and data to business premises, where copper-enhanced (xDSL), or Fibre-to-the-Curb (FTTC), architectures are in place.

Alternate Carriers or Value Added Network Suppliers

These companies are typically taking up licences to provide communications services in competition with the incumbent national operators. They will not own their transmission infrastructure. They must therefore, either buy bandwidth from the primary operator or form alliances or joint ventures with companies who have rights-of-way, or have deployed bandwidth.

In these cases cost and limited availability of bandwidth demand ATM's efficiency, and integration of voice and data services.

We should also note that many cellular operators are in this same situation. They must build a fixed voice network to interconnect their cell sites and message switching centres, and the efficiencies ATM can provide will result in major improvements in the cost-performance of their network.

Enterprise (Private) Networks

Most private, enterprise networks are buying bandwidth at commercial (retail) rates and must achieve the most they can with the resources on hand. In any corporation, a significant proportion of the traffic in the network is voice, so integration of voice and data under ATM becomes an obvious goal.

In many cases, such organisations will have already deployed a TDM network utilising E1 or T1 links. They will be looking to integrate these solutions into a new ATM network, and gain improvements in network performance and efficiency by moving from TDM to statistical multiplexing.

Two Basic Models Emerge

Given this perspective, two fundamental models for the transport of voice emerge. These are known as "voice trunking", and "voice switching".

Voice Trunking - This typically involved the tunnelling of voice traffic across a network between two fixed end points.

This is appropriate mechanism for the connecting of voice switch sites, PBXs, or message switching centres. There is no requirement on the network to be able to process or terminate signalling, other than the opportunity to use the signalling to detect idle channels.

However, the large scale of these networks often allows very efficient configurations to be achieved just by the analysis and use of traditional PSTN traffic engineering mechanisms.

Voice Switching - This involves the ATM network in the interpretation of voice signalling information and the routing of a call across the network.

The ATM switch, which receives the voice call, will route it to the appropriate destination. This type of functionality is most appropriate for VPN network.

Clearly, for this type of network solution to operate the ATM network must be able to interpret the signalling provided from the voice network.

In the past this has represented a major challenge, as many types of voice signalling existed, many of these being proprietary.

Today, these problems are being overcome. Many vendors provide ATM based solutions that are able to interpret the signalling provided by their own voice switches. More significantly, the widespread adoption of ISDN and QSIG voice signalling standards is allowing ATM vendors to offer a standardised voice-signalling interface.

The standards for interworking of narrowband and broadband ISDN signalling are also well advanced.

Network Requirements

What must the network do?

From this analysis, we can see a common set of network requirements emerge. To implement Voice trunking we must support the following minimum characteristics:

Characteristic	Necessity
Adaption	A mechanism to encode voice samples into ATM while meeting the delay and real-time constraints of voice traffic
Signalling	A mechanism to allow the end to end transport of voice signalling (Common Channel or Channel Associated) with the voice traffic.
Low cross network delay (Latency).	To minimise delay issues, and allow normal interactive conversation (Note this is not a requirement for broadcast applications)
Limited variation in delay.	To minimise delays and allow effective echo cancellation.

Beyond these basic demands, is a more extensive set of requirements necessary to support a complete voice switched solution, or implement efficient statistical multiplexing in a trunked voice environment.

Requirement	Necessity
Signalling Analysis	To allow set up and tear-down of circuits on demand (or allocation and release of resources)
Call switching and routing mechanisms.	To allow configuration of "real-world" VPN applications.
Silence Suppression or VBR encoding.	To realise statistical gain (provides at least a doubling in performance)
Call Admission Control (CAC)	To ensure quality of service is preserved
Network resource allocation	To allow statistical overbooking of network resources

The Objectives – What must be Achieved

Expectations

Any proposed ATM solution will be measured against the current generation of TDM solutions that have been deployed.

The ATM Forum has developed a set of solutions that are able to offer direct commercial or operational benefits to any user.

These allow voice traffic to be carried over an ATM network more efficiently than when carried over traditional TDM or packet based network infrastructures.

ATM's ability to statistically allocate network resources, and to accept VBR voice result in much higher levels of resource utilisation to be achieved. Idle channels do not consume bandwidth and resources can be allocated on-demand to voice and data traffic presented to the network.

Moreover, ATM's ability to transport real-time traffic over heavily loaded links results in an ATM solution being able to maintain call quality under all circumstances.

Two characteristics are the measure of success of any ATM based solution.

- Voice traffic must arrive unimpaired by its transport across the network.
- Valuable network bandwidth must be managed to provide a cost-efficient solution.

When we add ATM's ability to deliver multiple types of traffic over a single link, ATM becomes an overwhelmingly attractive solution.

How is it Done – What Solutions Exist?

The ATM Forum has defined three principal approaches to carrying voice over an ATM network. These are:

- CES (Circuit Emulation Service) — to carry full or fractional rate E1/T1 circuits between endpoints.
- Dynamic Circuit Emulation Service - DBCES
- ATM Trunking of Narrowband Services using AAL2 (under development)

A fourth approach is the transport of voice traffic that has previously been encapsulated in another protocol, for example Voice over IP or Voice over Frame Relay.

Each of these mechanisms exhibit challenges, benefits and problems, let us consider each in turn.

Circuit Emulation - CES

Circuit Emulation version was defined by the ATM forum in January 1997 as af-vtoa-0078.000. Today it represents a stable and reliable standard, which has been widely implemented by ATM equipment suppliers.

When using Circuit Emulation the ATM network simply provides a transparent transport mechanism for structured G.703/4 links. Voice is encoded into these links as in a normal TDM network using PCM, ADPCM, or other encoding & compression mechanisms.

The network will ensure that the delivered circuit is reconstructed exactly as received. CES is a full duplex mechanism, and presents the voice equipment with an apparent leased circuit. This approach is very valuable as no change to an existing, TDM or PBX network is required. A circuit-emulated link can in fact carry any type or mixture of traffic.

Circuit emulation uses the ATM AAL1 adaption mechanism to segment the incoming E1 or T1 traffic into ATM cells with the necessary timing information to ensure that the circuit can be correctly reassembled at the destination.

CES's advantages are the simplicity of implementation; the ATM network is used to provide virtual replacements for physical links in an existing network.

CES provides an ideal stepping-stone from legacy TDM networks to full ATM-enabled broadband solutions.

However, in its simplest form CES exhibits two limitations.

Firstly, it is unable to provide any statistical multiplexing. The ATM network does not differentiate between idle and active timeslots, this means that idle traffic/time-slots are carried.

Therefore, simple CES voice transport consumes about 10% more bandwidth than would be required to transfer the same voice traffic over leased circuits.

Secondly, it is often implemented as a point-to-point service - providing the transport of the contents of one network physical interface to a single other physical network interface. This can prevent the implementation of some network topologies, and can result in increased network cost, as a physical interface must be provided for traffic destined to each remote destination.

Dynamic Bandwidth Circuit Emulation - DBCES

The restrictions of simple CES resulted in the development of a new standard (Dynamic Bandwidth Circuit Emulation - DBCES) by the ATM Forum's membership.

This standard was ratified in July 1997 as af-vtoa-0085.000, and is implemented by many member companies in their equipment.

The objective of this standard is to enable dynamic bandwidth utilisation by detecting which time slots of a TDM trunk are active and which are inactive.

When an inactive state is detected in a specific time slot, the time slot is dropped from the next ATM Circuit Emulation data structure and the bandwidth it was using may be reutilised for other services.

DBCES may use any method of time slot activity detection. The specific implementation and method(s) chosen by individual vendors for activity detection is not defined, and various companies will adopt differing strategies.

The most common implementation mechanisms are the monitoring of A/B (on-hook/off-hook) bits in the channel associated signalling and the detection of idle codes within the payload of the voice channel.

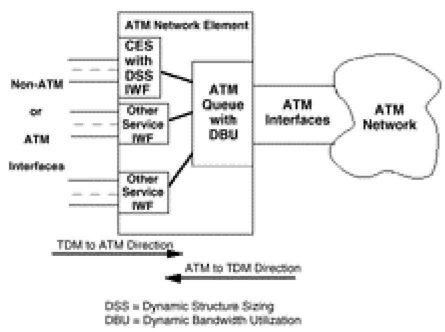

DBCES can operate in either PVC or SVC ATM network configurations. The active time slots are transmitted using the standardised CES service.

Figure 1.1

Figure 1.1 shows an internal picture of a DBCES equipped node. It shows the circuit emulation interworking function (CES IWF). This interworking function performs the following functions:

- Circuit Emulation Services.
- Time slot activity detection.
- Dynamic Structure Sizing (DSS) of the AAL1 structure which correlates with the active time slots in the TDM to ATM direction.
- Recovering the active time slots from the AAL1 structure, in the ATM to TDM direction, and placing them in the proper slots in the TDM stream.
- Placing the proper signals (e.g., ABCD) in each of the time slots of the recovered TDM stream.

In operation, the system assigns sufficient bandwidth to support the DBCES function when all the provisioned time slots are active.

When some of the time slots become in-active, the system dynamically stops transmitting the inactive time slots, thus fewer cells are queued for transmission.

The queuing system in the switch can then take the bandwidth not used by the DBCES function and temporarily assign it to another service. This capability provides bandwidth for UBR type services during times of lighter voice load. This increases the effective bandwidth utilisation of the network.

VBR Voice solutions using AAL2 Adaption

The forgoing CES mechanisms treat voice as being a constant stream of information encoded as a constant bit rate (CBR) stream. In reality voice is not like this at all, there are silences in conversation where one party speaks and the other listens.

There is no need to occupy bandwidth transmitting this silence.

Moreover, these mechanisms typically minimise the problems of cell construction delay by transmitting the voice as an uncompressed 64Kbps stream. This approach denies the network operator the opportunity to profit from voice compression technologies.

To address these two limitations the ATM forum defined a more advanced mechanism for the transport of voice as a variable bit rate compressed stream. This mechanism is described in the specification af-vtoa-0113.000, completed in February 1999, entitled "ATM Trunking using AAL2 for Narrowband Services".

This approach combines the suppression of silence in the conversation, with compression, and the ability to multiplex multiple voice channels into a single VCC. This multiplexing overcomes the packetization delay issues resulting from the use of low-bit-rate voice encoding.

This standard also optionally provides for the inclusion of network switching mechanisms based on the interpretation of the voice signalling channels. This will allow the building of switched private or public voice networks.

The AAL-2 adaption layer provides for the real-time delivery of variable bit rate traffic in an ATM network. This is perfectly matched to the needs of VBR compressed voice traffic.

The standard describes how data is collected from several TDM interfaces and compressed using one of a selection of compression algorithms. The resulting data streams are then merged into a single sequence of cells so that they can be transmitted over a single ATM virtual circuit. This results in each ATM cell potentially containing the data from several calls. Consequently, the packetization delay is much reduced, as the data rate perceived by the SAR function is the sum of the data rates of all the calls included in the multiplexed stream.

The interworking function defined in this standard may also route incoming traffic into various virtual circuits based upon a variety of parameters. This ability allows calls to be routed on called-address, incoming interface, time slot, priority or other signalling mechanisms.

For equipment which supports these "switched trunking" mechanisms, fully routed VPN networks can be built and significant reductions in network complexity can be achieved.

The benefit of Switching

In a point to point network, it is necessary to provide a virtual circuit (and sometimes physical port) for each destination in the network with which we wish to communicate. This results in an "N-squared" scaling problem, where network complexity, and number of VCCs required, grows geometrically with network size, as shown below.

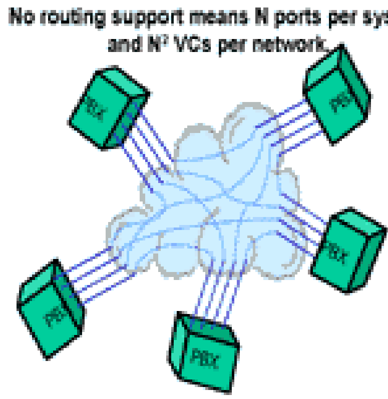

If the network can route calls by interpretation of the signalling, we can dramatically simplify the network design, and reduce the number of access ports.

In addition, if the access lines are being purchased this cost saving will be even greater. Moreover, the statistics of network provisioning will mean that more traffic per access port can be handled.

Interpretation of call signalling also allows the network designer to utilise the switching capability of the underlying network to avoid the need for "N squared" virtual circuits. The following picture shows this in practice.

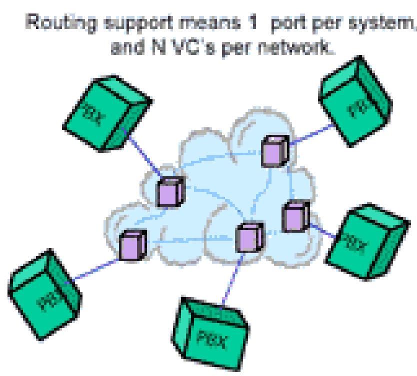

Clearly the benefits from this kind of solution are very significant and grow with network scale.

Above approximately 20 nodes, it is very costly to develop a workable architecture without some form of voice switching and routing.

User Benefits

A major benefit of ATM trunking using AAL2 for narrowband services is bandwidth savings. It can be achieved in these ways:

- By compressing voice — bandwidth allocation is less per call.
- By releasing bandwidth when the voice application does not need it — when the talker is silent or when call is completed.
- By routing and switching narrowband calls on a per call basis, further improvements in performance and efficiency can be achieved.

Encapsulation Techniques

The fourth mechanism available to the voice network architect who wishes to use ATM is to encode the voice traffic in another data protocol and then transport this protocol over the ATM network.

For example, the Voice over IP, or Voice over Frame Relay, initiatives which are gaining in strength can be supported when an ATM network is used to carry the resulting IP or Frame relay traffic.

This not really a "voice over ATM" solution, it should be viewed as the support of real-time data services.

Again, this is a simple implementation in as much as the network is being configured to carry simple data traffic. There are of course constraints on delay which must be met by the ATM network, but these are no worse than the criteria which must be met to enable voice over CES to operate.

The second benefit of this approach is that it is a statistical VBR mechanism. When the conversation stops most Voice FRADS or Voice over IP devices will stop generating frames or packets. This directly translates into savings at the ATM layer, as no cells will be generated for transport during periods of silence.

However, architecturally this is a complex solution where the ATM network is unaware that voice is being transported. Consequently the network is also unaware of the actual application requirements and therefore there is some inefficiency in the system.

For example the overlay frame or IP network will negotiate virtual connections between its nodes to cover its peak traffic demands; the ATM network must allocate resources to these virtual links.

Because the ATM network is unaware of the voice signalling, it is unable to release resources as calls terminate. The network can only assume that a short period of silence is in progress and therefore cannot reallocate its capacity to other applications, except on a "best efforts" or UBR basis.

This will lead to under utilisation of the ATM network as the CAC (Call Acceptance Control) applications are deprived of information on the actual state of the voice network.

A Comparison of Approaches

Table 1 summarises the benefits of the different VTOA trunking specifications; it shows the capabilities of the new standards defined in this area during the last 2 years. It is clear that the AAL2-based mechanisms are able to address all the principal areas of efficient voice transport.

Table 1

	Voice Compression	Silence Removal	Idle Channel Suppression	Switched Concentration
CES	-	-	-	-
DB-CES	-	-	÷	-
ATM trunking using AAL1 for narrowband services	-	-	÷	÷
ATM trunking using AAL2 for narrowband services	÷	÷	÷	÷
VoIP or VoFR over ATM	÷	÷	-	-

Conclusions

The bottom line

ATM is capable of transporting voice in an efficient and flexible manner.

There are a number of approaches available to the network designer. Some of these are well established and have been proven over several years, while other new approaches offer significant sophistication, and provide a highly efficient voice transportation mechanisms.

The range of standards now available allows a variety of applications and topologies of voice networking to be effectively addressed.

By including ATM in the design of the modern multi-service network, it is possible to integrate a complete range of voice and data services into a single network. This can translate into major savings in bandwidth and network complexity with the associated reductions in operational costs and improvements in reliability.

The vendors of the ATM Forum have pooled their knowledge to provide solutions, which are both robust and flexible.

So, consider speaking clearly with ATM in your network, and experience the benefits that will position you and your network for success!