DIGITAL AUDIO SYSTEMS – BEYOND THE CAT5 CONNECTOR

David Howe Shuttlesound (Electro-Voice), UK

1 INTRODUCTION

The trend to manage building systems over computer networks is only set to accelerate in the coming years as more and more manufacturers from all market sectors strive to add 'networking capability' to their products. The benefits to the client, in terms of infrastructure costs, operationally flexibility and future proofing an installations cabling plant can be significant if planned correctly but to capitalize on these advantages often requires a new skill set for the integrators and designers – Ethernet.

The audio professional, just like those working in many other building management industries, is having to meet the challenges of adopting complex, rapidly changing networking technologies within their designs. To add to the burden quite a high percentage of systems using networks for their digital audio transport will also need to incorporate at least some degree of emergency evacuation requirement. Resilient networks with multiple levels of hardware redundancy are typically required, so is real-time deterministic performance if the audio going in is to be accurately reproduced over the loudspeakers the network is to serve. Even if the audio consultant/engineer does not have to assume responsibility for the design and implementation of the network they will still undoubtedly have to take some responsibility for ensuring their audio system will not be compromised by the IT infrastructure.

This paper examines current networking technologies from an audio perspective and addresses some of the key issues and requirements of a network to ensure its successful implementation as an audio distribution medium.

2 NETWORKING BASICS

2.1 Ethernet: A Brief History

Ethernet is just one of several logical networking topologies. Logical topologies define the specifications for how data is transmitted across the physical medium of cables and hardware that interconnect computers and other networking devices. There are other topologies in use such as Token Ring and Fame Relay but Ethernet is by far the most popular with millions of nodes installed world wide.

Early attempts to network data between computers suffered from the problem that there was no mechanism to prevent the interconnected devices from trying to send their data at all at the same time. As they were connected in a daisy-chained bus configuration along a single cable each time more than one device attempted to send data, the data would become mixed up in a totally unpredictable way rendering it useless.

Ethernet, invented in 1973 by Bob Metcalfe, the founder of 3Com overcame the these limitations by implementing a standard called IEEE 802.3 CSMA/CD or *Carrier Sense, Multiple Access/Collision Detection*. The first Ethernet networks used the same single cable daisy-chained bus configuration

to interconnect the computers. Known as 10Base-2 and 10Base-5 these first Ethernet physical topologies, based on co-axial cables, were not particularly robust. A broken or disconnected cable between two devices would bring down the whole network. Bus topologies have all but died out today but they did pave the way for more flexible physical connectivity by proving CSMA/CD concepts.

As all the devices, or nodes within a network segment were interconnected by a single wire the segment became known as a *collision domain*. When two or more nodes attempted to send data simultaneously a condition called a *collision* occurred which invariably resulted in data corruption. CSMA/CD is the mechanism used to control data flow to try and prevent collisions and recover from them should they occur.

When a node has a data packet to send *Carrier Sense* tells it to first listen in case another node is currently transmitting, if so it will wait until the transmission is complete before attempting to send the packet.

Multiple Access describes the fact that a number of nodes can transmit and receive data on a single cable. Transmission by one node is typically received by all other nodes in the same collision domain.

Collision Detection defines what to do if two or more nodes attempt to send their data at the same time.

There is no ‘master’ controller or priority structure to regulate data flow, so each device has to operate independently. If the network is quiet and providing no other devices try to transmit, a node can send its data packet safely on its way. However, if another device does try to send its data at the same time a collision occurs. Both devices sense the collision and immediately cease transmitting, they then wait a random number on milliseconds for another quiet period and try again. Usually this will solve the collision problem. (See figure 2.1)

This arrangement worked well in the early days of Ethernet when the amount of data sent between computers was relatively small. Repeater hubs, described in the next section, overcame the inadequacies of the physical bus topologies and small reasonably reliable networks could be constructed. However, Ethernet networks implementing CSMA/CD have two major shortcomings;

1. CSMA/CD breaks down under pressure. In a heavily loaded network carrying large amounts of data the contention for bandwidth between nodes can cause the network to become so overloaded that it crashes. Nodes will time-out because they cannot send their data in an acceptable timeframe and eventually the whole dataflow will become so disjointed that hardly any of the devices can operate correctly. The larger the collision domain the more likely it is that collisions will occur. Or, put quite simply CSMA/CD just doesn’t scale very well.
2. Data is transmitted onto the network as electrical pulses so it will travel at high speeds but still requires a finite amount of time to travel to the furthest most nodes. In the case of the two most distant nodes in a collision domain the maximum length of the cables interconnecting these two nodes, also known as the *Network Diameter* has to be limited in size. There are the normal losses and attenuations caused by the cables themselves but more importantly if both of the end nodes attempt to send data at the same time small data packets could collide in the middle of the cable and become corrupted. The nodes would not realise that a collision had occurred because the transmission was complete before the collision took place. To overcome this situation the minimum data packet size has to be large enough to keep a node in transmit mode long enough for a collision to propagate back from the furthest node in the collision domain. Ethernet standards set a minimum data packet size of 512bits, with smaller packets being automatically padded to this size. For a Fast Ethernet 100Base-T repeater hub based network the network diameter must still be limited to 200 meters to ensure reliable collision detection.

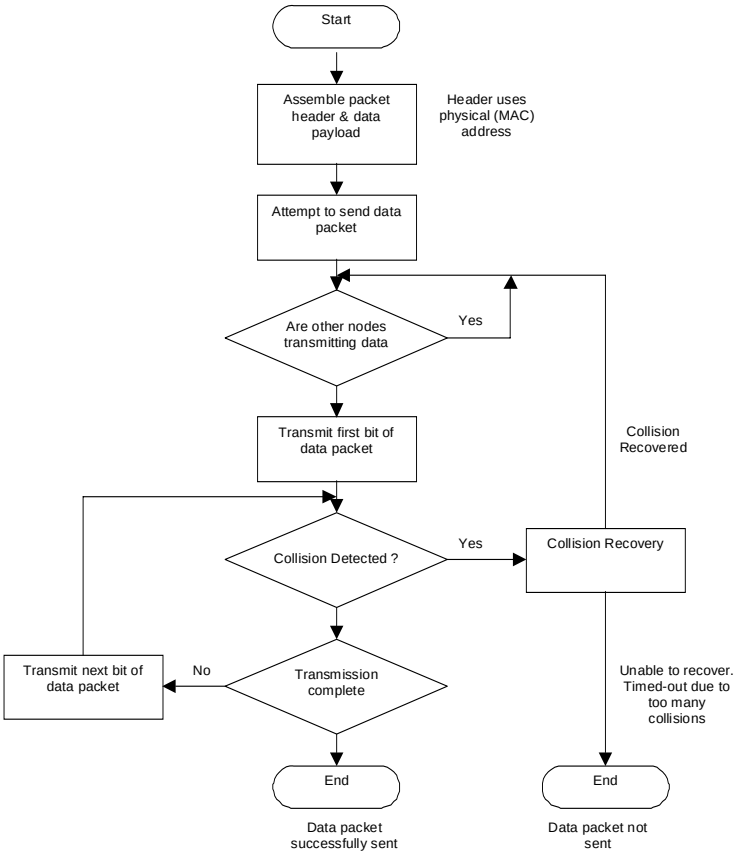

Figure 2.1 shows a simplified flow diagram of how Collision Detect works.

2.2 The OSI Model

During the 1980's a working group attempted to create a logical model for the various parts that make up a network, the model was called *Open Systems Interconnection* (OSI). It was used to define modes of interconnection between different components in a networking system allowing the physical method of transport to be designed independently to the protocols and applications running over it. When we talk about 'Layer 2' and 'Layer 3' networking it is the layers of the OSI model that we refer to. The model is not complicated but as the layer numbers increase so does the level of abstraction. (See figure 2.2)

The Application Layer (Layer 7)
The top layer in the stack, the Application layer is where the end-user application resides. Many protocols are defined for use at the Application layer, such as HTTP, FTP, SMTP, and Telnet.

The Presentation Layer (Layer 6)
The Presentation layer is used to provide a common way for applications (residing at the Application layer) to translate between data formats or perform encryption and decryption. Mechanisms to convert between text formats such as ASCII and Unicode may be considered part of the Presentation layer, along with compression techniques for image files such as GIF and JPEG.

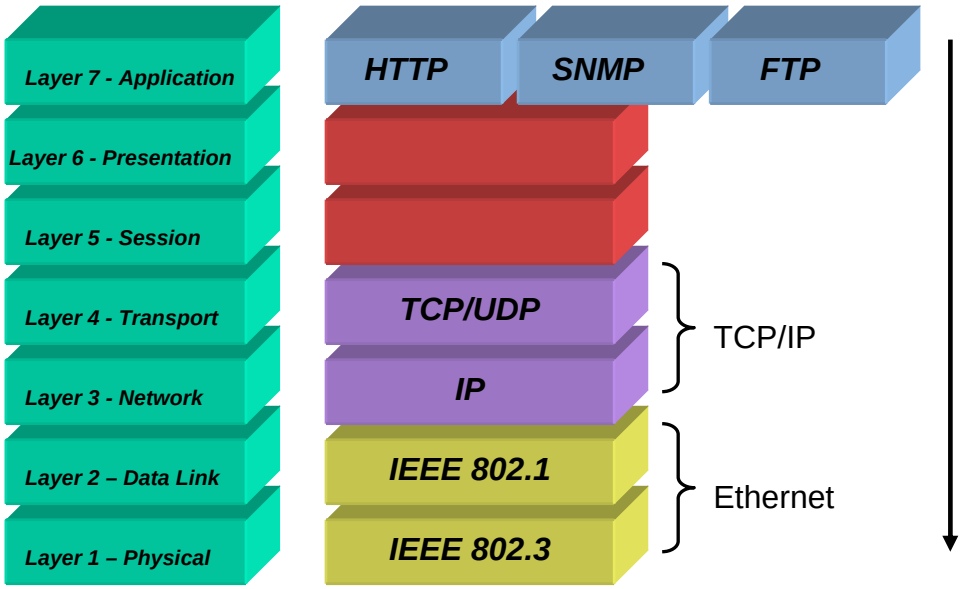

Figure 2.2 The OSI Model

The Session Layer (Layer 5)
The Session layer handles the order of data packets and bi-directional communications by coordinating multiple Presentation layer processes communicating between end devices. The Session layer is used by applications at either end of the communication to tie together multiple Transport layer sessions and provide synchronization between them. For example HTTP protocol can use multiple TCP connections to retrieve objects that make up a Web page. The Session layer provides application coordination between these separate TCP connections.

The Transport Layer (Layer 4)
The Transport layer handles the transport mechanism between devices. It is concerned with ensuring data packets get to their destination intact. Error checking and delivery failure notifications are processed by the transport layer.

The Transport layer is the first at which we see the concept of packets or datagrams of data that will be transported across the network. TCP and UDP are examples of Layer 4 protocols used to provide a delivery mechanism between devices.

The Network Layer (Layer 3)
Whereas Layer 4 is concerned with transport of the packets within a communication channel, the Network layer is concerned with the delivery of the packets. This layer defines the addressing structure of the network and how packets should be routed between devices. The Network layer typically provides information about which Transport layer protocol is being used, as well as local checksums to ensure data integrity. Internet Protocol (IP) and Internet Packet Exchange (IPX) are examples of Network layer protocols.

Traditional Internet routers operate at the Network layer by examining Layer 3 addressing information before making a decision on where a packet should be forwarded. Hardware-based Layer 3 switches also use Layer 3 information in forwarding decisions. Layer 3 routers and switches are not concerned with the type of data the packet contains, but simply where the packet is flowing to and from.

The Data Link Layer (Layer 2)
The Data Link layer also defines a lower level addressing structure to be used between end systems as well as the lower level framing and checksums being used to transmit onto the physical

medium. Ethernet, Token Ring, and Frame Relay are all examples of Data Link layer or Layer 2 protocols.

Traditional Ethernet switches operate at the Data Link layer and are concerned with forwarding packets based on the Layer 2 MAC addressing scheme. Layer 2 Ethernet switches are only concerned with where the MAC address of the recipient device resides.

The Physical Layer (Layer 1)
As with all computer systems, networking is ultimately about making, moving, and storing 1s and 0s. In networking terms, the Physical layer defines how the user's application, or networked device data is turned into 1s and 0s to be transmitted onto the physical medium. The Physical layer defines the physical parts of the network such as cabling and interface specifications. 100Base-TX, Gigabit and RJ45 are all examples of Layer 1 specifications.

Interestingly the OSI model was never actually implemented as a network protocol; instead, the existing protocols, primarily TCP/IP, were refined using the OSI reference model.

2.3 Hubs and Switches

2.3.1 Hubs

A hub, full name repeater hub, is a piece of Ethernet hardware with multiple ports to interconnect various nodes. A data packet transmitted by a node is received at the hub which then simply re-clocks it, boosts the signal and sends it out to all ports except the port it was received on. If one of the ports is connected to a port of another hub the size of the network can be expanded in a star configuration as the second hub will also resend the data packet it receives to all of its ports except the receiver. Each time data passes through a hub in this fashion it is known as a *hop*.

As hubs simply resend any data they receive to all connected nodes on the network, whether the nodes require the data or not, the available bandwidth has to be shared across the entire network. Hubs also can only operate in half-duplex mode, they cannot receive and transmit data at the same time, therefore only one node is allowed to transmit data onto the network at any one time. If two nodes attempt to transmit their data at the same time a collision will occur resulting in lost data. Hub based networks are therefore part of a collision domain.

Due to the shared bandwidth and network diameter limitations, hub based networks have severe scalability limitations, especially for the high bandwidth requirements of multi-channel digital audio. Fortunately hubs are no longer used for constructing networks but their operation and limitations should at least be understood as many older networks may still rely on this technology. Under these circumstances it is highly recommended that rather than trying to integrate digital audio onto this network a new infrastructure is considered.

2.3.2 Switches

A switch, like the hub is a multi-port device used to interconnect multiple nodes. Unlike a hub the switch has intelligence. It is able to read the source and destination MAC addresses of each data packet and forward it only to the port on which the destination device resides. Data sent to a single destination like this is known as *Unicast* addressing. A switch is still able to forward data to all ports if necessary although this time, unlike the hub which will always send data to all of its ports, information contained in the data packet is used to determine the destination. When data is transmitted to all ports the term *Multicast* addressing is used. CobraNet uses the same terminology to differentiate between point to point and point to many point audio distribution. Figure 2.3 illustrates the concept of Unicast and Multicast addressing.

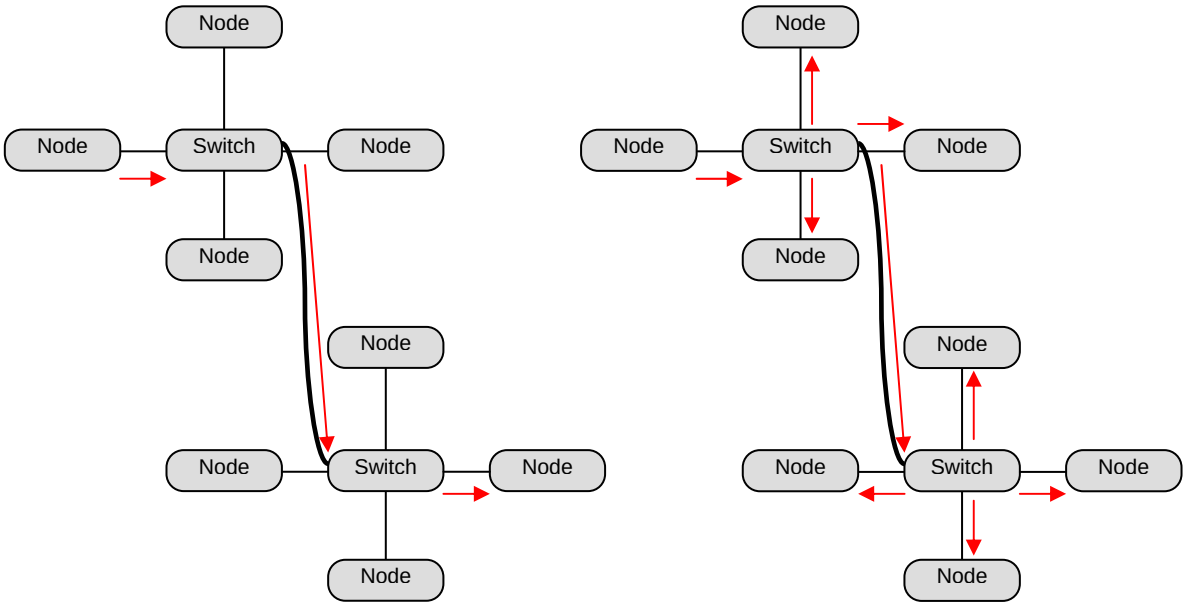

Figure 2.3 Shows the differences between Unicast and Multicast Addressing

Each port on the switch operates in full duplex mode, i.e. it can send and receive data simultaneously. This is achieved using two twisted pairs of a CAT-5 cable, one pair for transmit and the one pair for receive. The switch fabric manages data throughput by queuing data packets in its internal FIFO (First In, First Out) memory buffer. This process is known as *Store and Forward*, where data packets are stored before the switch reads the destination MAC address and then forwarded to the correct destination port. The concept avoids placing unnecessary traffic on other ports and coupled with the fact that multiple ports can transfer data simultaneously, significantly increases available bandwidth. As the switch actively controls data flow CSMA/CD is no longer required, thus removing the network diameter limitations of hub based networks. Therefore, much larger and efficient networks can be constructed.

Initially, when first powered up a switch has to ‘learn’ the destination addresses of devices attached to it. This is the MAC (Media Access Control) address of each devices network interface. Every piece of networking hardware has its own unique IEEE issued six-byte long MAC address which could be considered as the low level serial number of the device. Higher level protocols such as TCP/IP addressing structures map IP addresses to the physical MAC address of a device so although an IP address may be used in software applications to communicate with a device, in the background data packets are still routed using the MAC address on Layer 2 switches.

As nodes start to send data packets the switch reads the source MAC address and stores it in a look up table so it will know where to send data addressed to that particular device. If the packet contains a destination address not currently stored in the look up table the switch will ‘flood’ the packet out of all ports in a bid to find the destination device. If the device is connected to the network it acknowledges receipt of the packet and its MAC address is also stored in the lookup table for future use. After a short learning period the switch will know the destination port of all connected devices making routing of data packets quite an efficient operation.

Similarly if no packets are received from a device for a certain amount of time (typically a few minutes) it is removed from the look up table. In this way as devices are connected and disconnected the switch is able to relearn new routes to destinations. Fortunately too, the whole idea of the complete network crashing because a node has been disconnected is a dim and distant memory.

If switches in a network are connected in a loop a *Broadcast Storm* will result where a single packet will circulate endlessly. This condition quickly consumes all available bandwidth on the loop making the network unusable. Rapid Spanning Tree Protocol discussed in Section 5 is used to prevent this situation where loops are a desirable feature from a fault tolerance point of view.

Switches will also add latency to a data packet which is directly proportional to the packet size. Store and Forward adds a minimum latency of one packet time per switch 'hop' ranging from 5us to 120us per switch depending on the packet size. Switch processing adds latency in the order of 5us per switch while management protocols such as QoS (Quality of Service) can deliberately hold back low priority packets on a busy network while higher priority packets get pushed through. Latency must be considered if the forwarding time of exceptionally time sensitive audio data packets is to remain within the required specifications.

Unmanaged and Managed Switches
Apart from the fact that switches come in various frame sizes with differing numbers of ports and medium types, CAT5/6 RJ45 connectors, fibreoptic connectors etc. they are also available as either managed or unmanaged variants.

Unmanaged switches are typically simple 'plug and play' devices with no configuration options. They are very cost effective and ideal for small, simple networks which don't require any redundancy. A simple unmanaged switch network can be used for both CobraNet and EtherSound networks but in this situation it is recommended that the network is dedicated entirely to the audio distribution as there is no way to manage and control the priority of other traffic. If the network is heavily utilized it is possible that indiscriminate loss of packet data may occur as bandwidth is exhausted bringing with it audible 'glitches' or dropouts in the audio signal.

Managed switches on the other hand, whilst performing the same basic function as the unmanaged type come with a range of configuration options to better control data flow through the network. They cost more than the unmanaged units but in all but the simplest networks the additional expense is well justified. It's also quite feasible to incorporate both managed and unmanaged switches in the same network if careful consideration is given to the operational requirements of the design.

Typically managed switches will provide some of the following additional configuration features, choice of vendor and model can be based on the configuration options required for the project;

- User interface via RS-232, Telnet, SNMP or web browser allows configuration options to be assigned and saved. Many of these user interfaces can also be used for diagnostic and management duties providing the network administrator with information on the switches data throughput and potential fault and failure information.
- RSTP (IEEE 802.1w) *Rapid Spanning Tree Protocol* for fault tolerant loop architectures
- VLANS (IEEE 802.1Q) *Virtual Local Area Networks*
- QoS (IEEE 802.1p) *Quality of Service*
- Link Aggregation (IEEE 802.3ad)
- Port Mirroring and more ……

These are the most useful configuration features of a managed switch when designing a network for distributing digital audio. RSTP will be discussed in later but a brief summary of the other features is warranted

Virtual Local Area Networks (VLANS) is a mechanism to provide multiple logical network segments or subnetworks on the same physical network infrastructure. A VLAN subdivides the switch into two or more logical switches with separate broadcast domains. As such data sent by a node can only be received by other nodes in the same VLAN group. If audio data has to co-exist with other data on the same network VLANS help to segregate the data and maintain the

deterministic requirements of the audio protocol specification. Bandwidth Reservation allows sufficient bandwidth to the allocated to the audio VLAN so data on other VLAN's will not cause glitches or dropout. Critical real-time audio traffic can be isolated from data which means audio devices won't have the processing overhead for unrelated traffic and from a security point of view VLAN's restrict traffic to the required devices. Devices connected on different VLAN's cannot usually interact.

Quality of Service (QoS) provides consistency to audio traffic data flow using bandwidth reservation. As the network becomes busy audio data will be pushed to the top of the switches Store and Forward queue resulting in predictable latency even when the network is highly utilized. Class of Service (CoS) provides a prioritization structure for data packets. In the case of audio data packets carrying emergency evacuation messages, correctly setting the CoS will ensure these data packets always get through no matter how busy the network.

Link Aggregation (also known as Trunking) can be used in larger networks where switches are linked together. Typically a single link will have the same bandwidth as all other ports on the switch even though it may be called upon to carry significantly more traffic. Link Aggregation allows multiple physical connections between switches and balances the load between the links. Under normal circumstances these multiple links would create a loop condition, Link Aggregation prevents this. It also provides a method of redundancy, if a link should fail the load will be balanced across remaining links to compensate. In order to use Link Aggregation ports have to be manually assigned as members of an aggregation group.

Port Mirroring is a diagnostic function that forwards a copy of each packet sent and received from one port of the switch to another where they can be monitored to keep track of switch performance and data throughput at various ports. It can be used to answer questions like, is the data packet actually getting to the device, when problems occur. Port Mirroring places a high overhead on the switch fabric and hence should be used sparingly and deactivated when not required.

Switch Speed

There are currently several different speeds at which Ethernet can operate

- 10 MBit/sec – 802.3 10Base-T Ethernet
- 100 MBit/sec – 100Base-T also known as Fast Ethernet
- 1000 MBit/sec – 1000Base-T Gigabit
- 10,000 MBit/sec – 10 Gigabit is on the way
- 11 MBit/sec – 802.11b wireless LAN
- 54 MBit/sec – 802.11g wireless LAN

Both CobraNet and EtherSound require 100MBit/sec network speed as a minimum. Because of this it is not possible to connect nodes wirelessly – at least not yet!

Cable Length Limitations

Fast Ethernet (100Mbps)

- 100Base-TX (Cat5/Cat5E) 100m
- 100BASE-FX (62.5 m Fibre) 2000m

Gigabit (1000Mbps)

- 1000Base-T (Cat5E/Cat6) 100m

- 1000BASE-SX (62.5 m Fibre) 220m
- 1000BASE-LX (62.5 m Fibre) 550m
- Greater distances are possible up to 70km using long haul single mode fibreoptics

3 COBRANET™

CobraNet is a combination of hardware (the CobraNet interface), network protocol and firmware. It can operate on a switched Ethernet network or a dedicated repeater hub network and provides the following communication services

- Isochronous data transport
- Sample clock distribution
- Control and monitoring data transport

The CobraNet interface performs synchronous to isochronous and isochronous to synchronous conversations as well as the data formatting required to transport real-time digital audio over the network. Since CobraNet is Ethernet based (it operates on Layer 2 of the OSI model), in most cases normal data packets and CobraNet digital audio data packets can co-exist on the same physical network.

CobraNet Terminology

Conductor – The conductor is the CobraNet interface elected to provide master clock and transmission arbitration for the network. All other CobraNet devices on the network are said to be operating in performer role. Clock distribution is via a beat packet transmission timestamp. Performers adjust their local clock so it is always in sync with received beat packets. The conductor also arbitrates bandwidth and network routing resources. If a CobraNet device wants to transmit onto the network, it puts out a reservation request to the conductor. The conductor approves the request by adding a permission entry to the beat packet. The request may be denied by the conductor if there is insufficient network bandwidth or if the requested bundle is already in use by another transmitter.

Bundle – A bundle is the smallest networked audio routing envelope. A bundle represents transmission of an Ethernet packet once per isochronous cycle. It may carry 0 to 8 audio channels. Each bundle is assigned a number between 1 and 65279. Each bundle can only have a single transmitter and there are two bundle type classifications;

Multicast Bundle – Bundles 1 through 255 are designated Multicast bundles. A transmitter configured to send multicast bundles will always transmit the audio regardless of whether any devices are set to receive it. Point to multipoint connections are possible. Any devices that wish to receive the multicast bundle can set their receivers to the same bundle number and receive the audio.

Unicast Bundles – Bundles 256 through 65279 are Unicast. Only point to point connections are allowed. A transmitter and receiver must both be set to the same bundle number before any data flows. There is also a MultiUnicast mode which allows a transmitter to send its bundle to up to four receivers although this is not a well documented feature and may not be implemented in all manufacturers products.

A 100Mbps network will support up to 64 audio channels on each port while using a Gigabit backbone allows several hundred audio channels to flow over the network.

There is one scalability limit that designers should be aware of when using CobraNet and that is the limit on multicast bundle traffic. As multicast bundles are forwarded to all CobraNet devices network wide only a limited number can be accommodated on the network if port saturation is to be avoided. The CobraNet specification recommends that under normal circumstances no more than four multicast bundles are used. If more are required it may be possible to use MultiUnicast or separate the network into separate VLAN's.

CobraNet has a guaranteed fixed latency of 5-1/3ms between CobraNet devices providing the network meets the specification requirements for forwarding delay. Lower latency down to 1-1/3ms can be achieved at the cost of higher bandwidth requirements.

4 ETHERSOUND™

EtherSound technology is also compatible with IEEE 802.3 standards and like CobraNet operates on switched, Fast Ethernet networks. It too operates on Layer 2 of the OSI model. However unlike CobraNet which uses the principle of Multicast and Unicast bundles EtherSound audio transmission is in one direction only with audio outputs only being available 'downstream' of a switch. Bi-directional audio transport is possible in daisy chain configuration on certain devices (see figure 4.1).

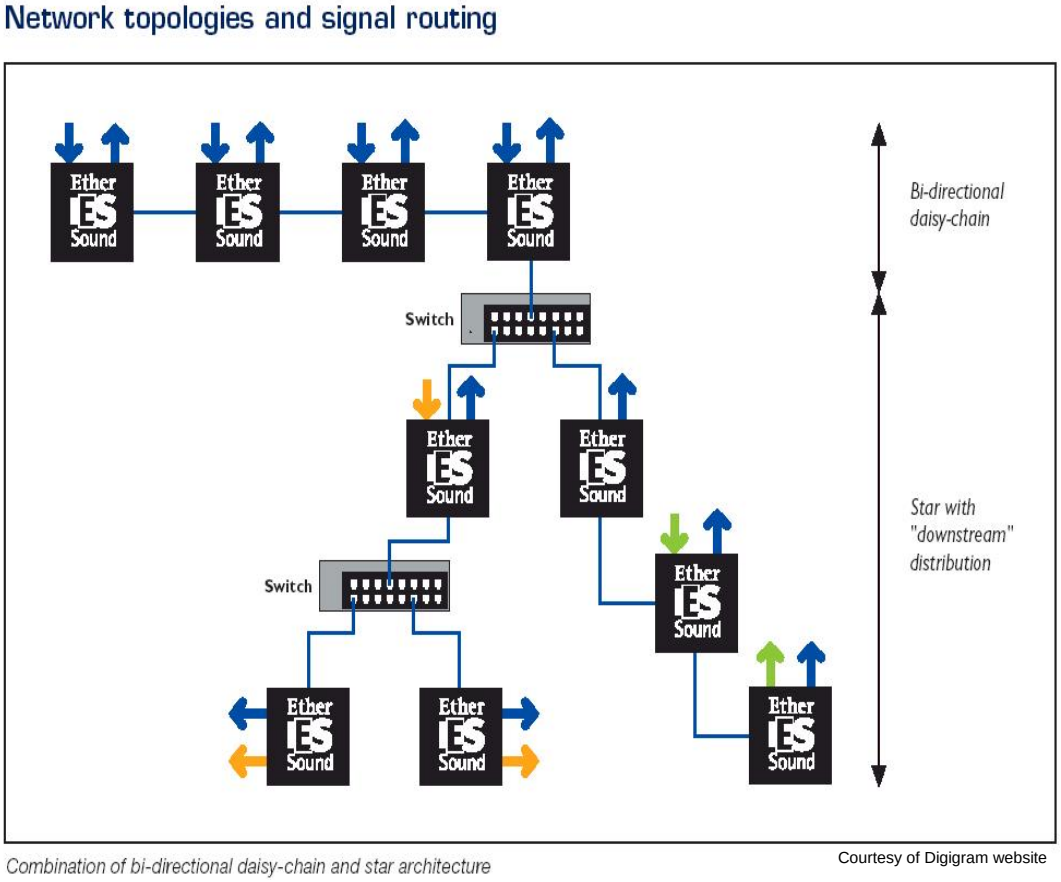

Figure 4.1 EtherSound Networking Topologies

Unmanaged switches may be used to distribute EtherSound transmissions to other EtherSound devices allowing more complex network architectures and extending distance between devices.

One key specification of EtherSound is its incredibly low latency of 1.22us per EtherSound device hop when passing audio downstream, however EtherSound does not use standard Ethernet data packet structures so if it is to be incorporated into an existing network this will only be possible by segregating EtherSound traffic with a VLAN.

EtherSound using a Gigabit backbone will be available shortly allowing some 256 audio channels to be transported across the network in each direction.

Both CobraNet and EtherSound are audio manufacturer independent protocols. This has the distinct advantage that the protocols have been implemented in a wide variety of varying manufacturers equipment and unlike many of the vendor specific protocols that are around they both use standard Ethernet networking hardware. It is quite feasible to interconnect CobraNet devices from several different manufacturers and be able to pass audio, the same holds true for EtherSound. Both protocols have their own set of advantages and pitfalls so the designer must consider these carefully when choosing the protocol for a particular project as unfortunately its not possible to get the best of both worlds and connect a CobraNet device to an EtherSound device and pass audio.

5 FAULT TOLERENT NETWORKS

RSTP (Rapid Spanning Tree Protocol) allows the creation of fault tolerant rings that incorporate redundant links which are blocked to prevent loops. The resulting tree (or star) configuration 'spans' all switches but eliminates loops, which are not allowed, as they result in a broadcast storm where data packets circulate endlessly rendering the network useless. It's a very similar situation to acoustic feedback and generally requires that the whole network be rebooted to recover.

RSTP is a refinement of the original Spanning Tree Protocol (STP) which has two primary parameters that must be configured for each switch on the network. They are; bridge priority and port priority. The switch with the lowest bridge priority becomes the root switch in the network. The root switch is the logical, although not necessarily physical, center of the of the network and may change over time as switches or links are removed or fail.

The goal of STP is to ensure that only one switch is responsible for forwarding traffic from the direction of the root switch onto any given link. If there is only one active path from the root to a link then there will be no logical loops in the topology.

The STP protocol does suffer from a number of drawbacks that limit its suitability for redundancy in audio networks, namely;

- STP has long failover and recovery times. When a link fails a backup path to the root switch can take upward of 30 seconds to recognize that it is the best (or only) path to the root and become usable.
- When a failed link returns to service, information about the 'better' path will instantly cause the backup link to start blocking, but the segment of the network below the link that is returning to service is isolated, again for upward of 30 seconds, until the link relearns its path to forward data.
- Another problem with STP is that all links must initially go through a lengthy period of address learning even if a link is point to point between nodes on the same switch. This has the knock on effect that bringing the network up after power up, or after a power outage can take several minutes.

RSTP solves the original STP protocols lengthy failover and recovery times when a physical link fails by a number of means. Whereas STP switches store only the best path to the root switch,

RSTP switches store all potential paths. When links fail RSTP already has pre-calculated routes to fall back on. Additionally, unlike STP switches, an RSTP switch will respond to another switch that advertises an inferior or incorrect route to the root switch allowing the switch with incorrect information to be rapidly retrained.

RSTP solves STP's problem with lengthy recovery times by introducing a new procedure called proposing-agreeing. Proposing and agreeing works after a better path to the root switch is restored by shuffling the restored part of the network one hop at a time towards the network edge. This method also enables the network to come up quickly on power up.

RSTP also introduces a method for quickly bringing up ports at the edge of the network while still protecting them against loops. If the port is designated as an 'edge' port (a port that typically connects a node to the network), RSTP will continue to send configuration messages known as BPDU's (Bridge Protocol Data Units) out of the port to detect loops but will allow data traffic to flow as soon as the port rises. Nodes connected to edge ports can send data traffic without the extensive delays imposed by STP.

Practical rings using STP were limited to seven switches. RSTP has increased that limit to 31 switches although some manufacturers have enhanced the specification still further and can support rings with up to 80 switches. Latency caused by excessive switch hops should still be considered for a deterministic audio network but using RSTP with Gigabit backbone links will allow considerably more that the seven switch limit imposed by STP.

As a simple rule of thumb 5-10ms of recovery time can be assumed for each switch in the ring, for example a network consisting of 10 switches would therefore recover in under 100ms in the event of a link failure. This might cause a minor glitch in the audio but would be pretty much imperceptible.

Figures 5.1 to 5.4 help illustrate several possible networking topologies for a CobraNet based audio network.

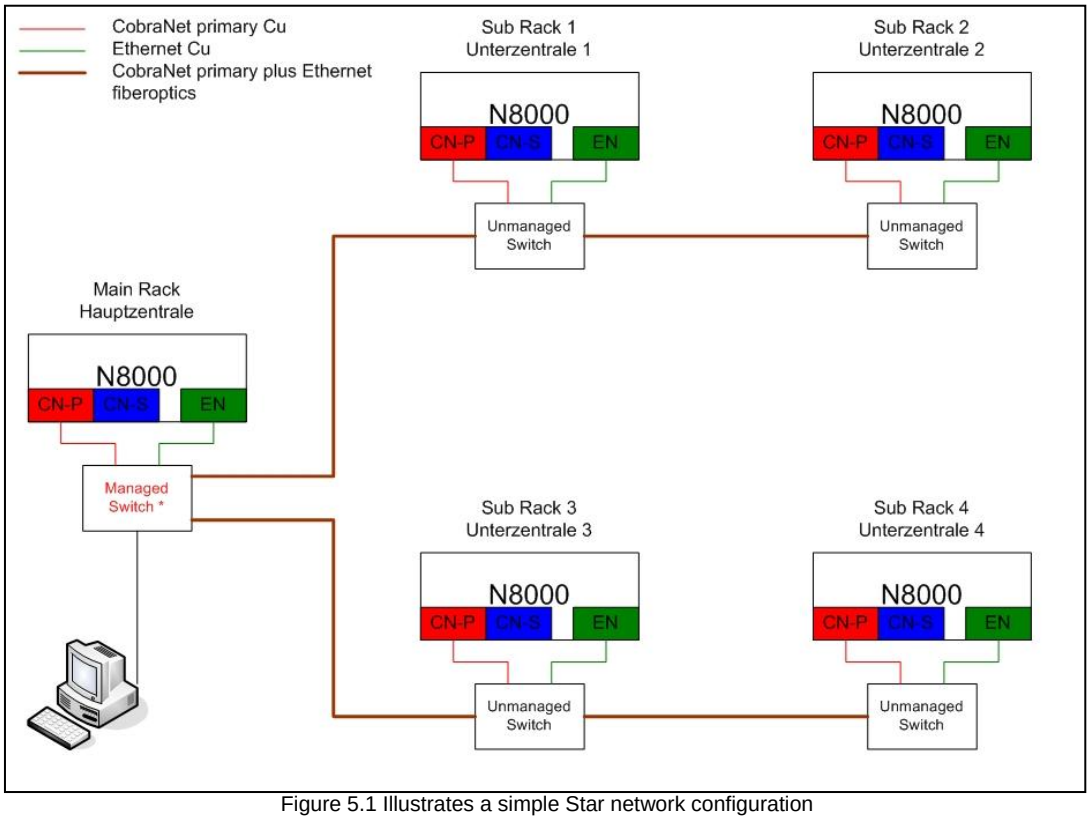

Figure 5.1 Illustrates a simple Star network configuration

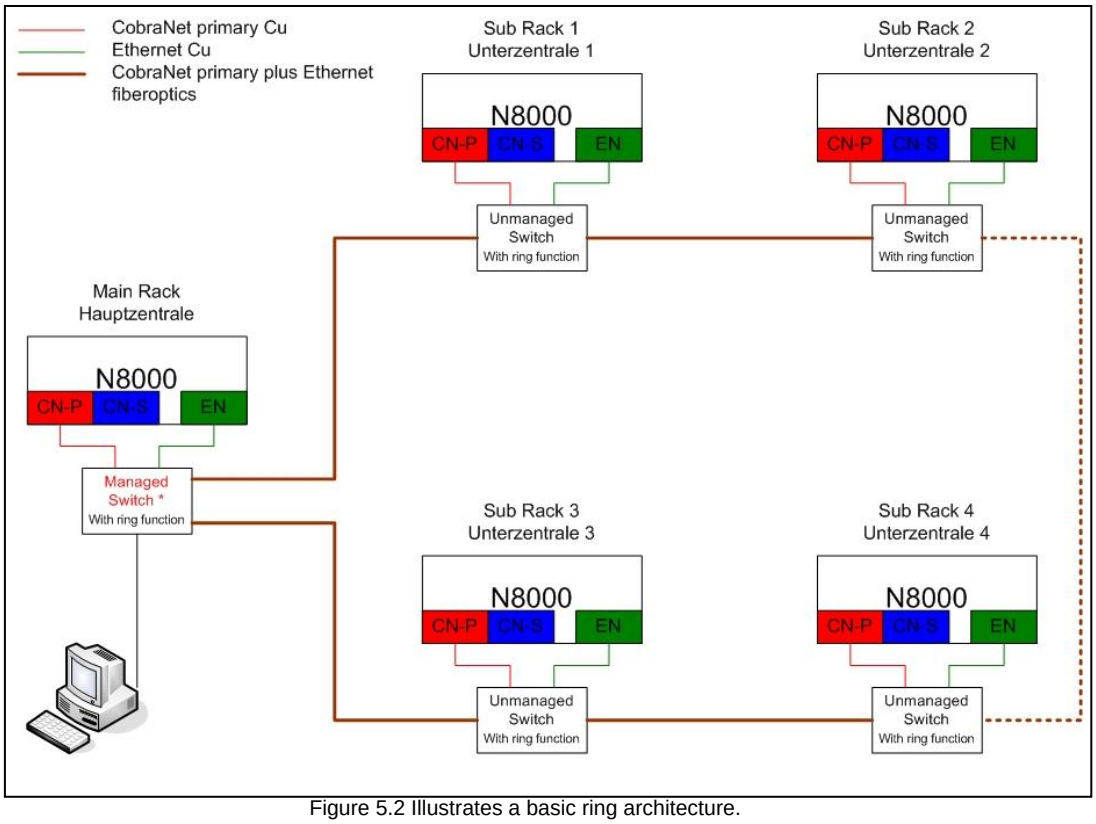

Figure 5.2 Illustrates a basic ring architecture.
The dotted line represents a physical cable link that might have a logical block created by RSTP implementation.

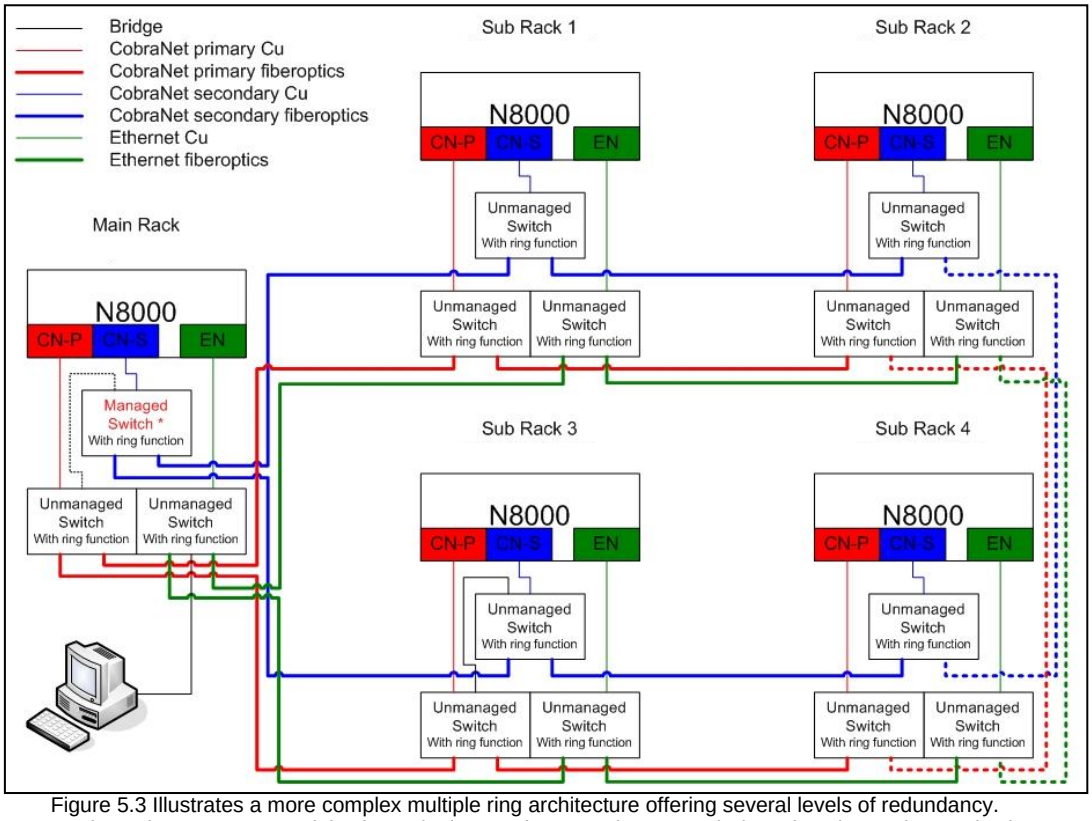

Figure 5.3 Illustrates a more complex multiple ring architecture offering several levels of redundancy.
Note the CobraNet CM-1 modules have dual network connections CN-P(Primary) and CN-S(Secondary)

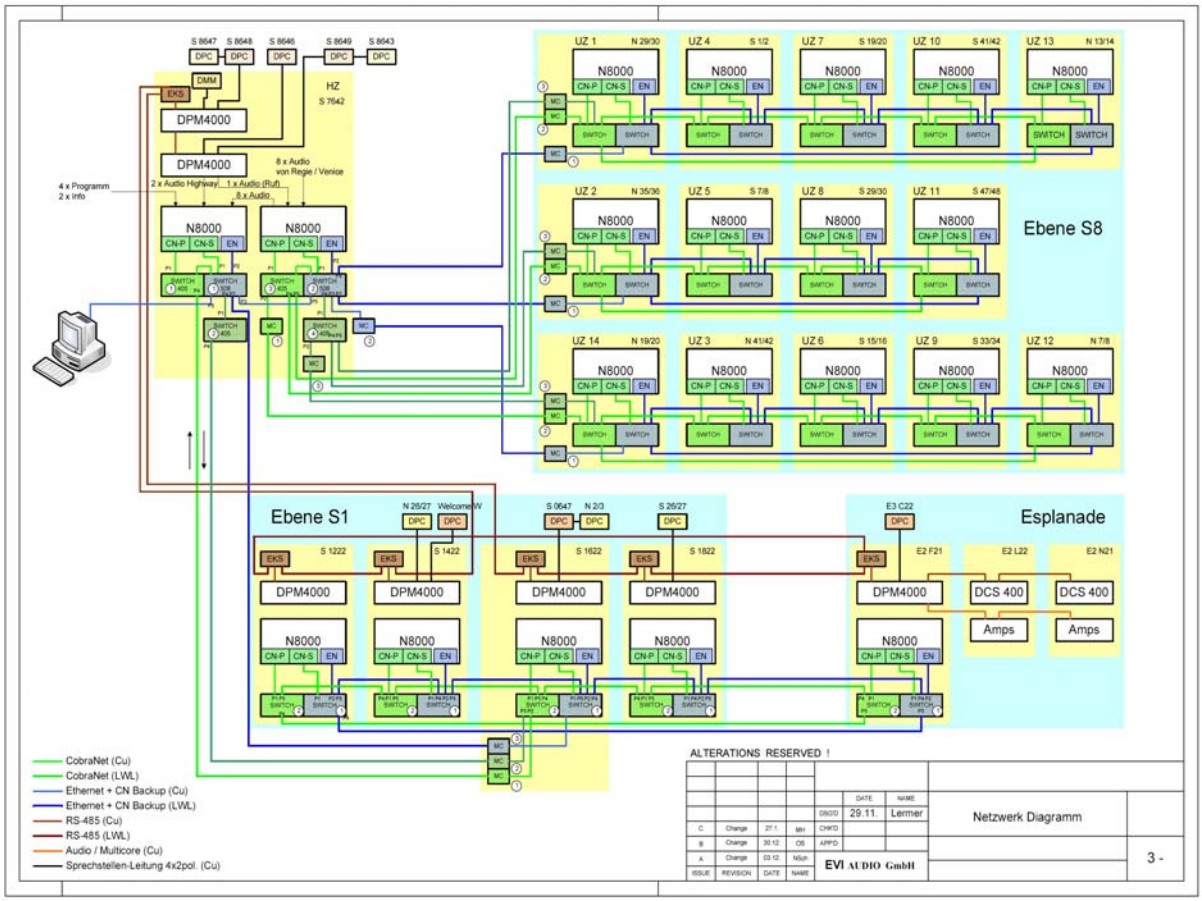

Figure 5.4 Illustrates a possible design for a real world digital audio network using CobraNet and the ElectroVoice NetMax N8000 digital audio engine.

6 SUMMARY

Networking is a huge topic and this paper has only really been able to scratch the surface, however key concepts have been introduced which will hopefully entice the reader to research further. Computerised networks are here to stay …. at least in our life times, and their usage will only continue to dominate. Audio engineers who understand the concepts have the tools at their disposal to create resilient, fault tolerant networks with more than sufficient bandwidth for their current demands, the challenge is to effectively plan so the network can adapt to future at present unknown demands.

Clients are looking to integrate more and more data and control throughput on their networking investments, this is not necessarily a bad thing but real-time audio distribution places overhead on the network that the client may not appreciate technically ….. and they will certainly not appreciate aurally if audio data packets start getting delayed or lost. Our requirements are not always understood by IT professionals either so it has to be our mission to impart our knowledge to them all so audio distribution can truly be brought into the 21st Century.

7 REFERENCES

1. ANSI/IEEE Std. 802.1p
2. ANSI/IEEE Std. 802.1Q
3. ANSI/IEEE Std. 802.1w
4. ANSI/IEEE Std. 802.3
5. The Switch Book, Rich Seifert, Wiley
6. www.cobranet.info
7. www.ethersound.com