

An introduction to networked audio

This white paper's subject is 'Networked Audio'.

Audio networking introduces new exciting possibilities for the professional audio industry. But it also drastically changes the way audio systems are designed, built and used, introducing new technologies and strategic issues to consider when investing in a networked audio system.

In this white paper the basics of audio networking will be covered in a straightforward comprehensive format. We assume the reader has an advanced knowledge of analog audio systems, a basic knowledge of digital audio systems and no knowledge of computer networking. This white paper is only a basic introduction to the subject; for detailed information we refer to the many documents on the internet made available by the IT equipment manufacturers around the world.

The Yamaha Commercial Audio team.



The complete package

An introduction to networked audio	
1.	What is networked audio?
2.	Three good things to know about networked audio
3.	Three things to take into consideration in networked audio
4.	What is an Ethernet network?
5.	Network topologies
6.	Redundancy concepts
7.	Cabling
8.	More about CobraNet™
9.	More about EtherSound™
10.	System engineering
11.	Investing in a networked audio system
12.	Networked audio glossary

1. What is networked audio?

Networked Audio

With the introduction of digital technologies the amount of information a single cable can carry has increased from a few thousand bytes in the sixties to a few billion bits in 2006. Regular affordable connections in every day information systems now carry one gigabit of information in a single fiber cable over distances spanning many kilometers.

This bandwidth is enough to transport hundreds of high quality audio channels, replacing hundreds of kilograms of cabling in conventional analog systems. More important, the functional connections in a networked audio system can be designed to be completely separated from the physical connections in the network. This functionality opens up a wide array of exciting possibilities for the audio industry: any number of i/o locations can connect to the network anywhere in the system without the limitations of bulky cables, leaving the actual connections to be managed with easy to use software. A networked audio system is digital so audio connections are kept in the digital domain, far away from electromagnetic interferences and cable capacitances degenerating analog audio quality. Control signals can be included in the network without additional cabling. Computers can use the network to control and monitor audio devices such as digital mixers and DSP engines. Video connections can be included using affordable IP cameras; and so forth.

Digital audio distribution

There are many systems on the market that distribute audio over a single cable using copper or fiber cabling supporting Point To Point connections such as MADI. These formats are called 'P2P' systems, they connect location A to location B, offering cost effective solutions for fixed designs with two locations.

EtherSound™

The EtherSound™ protocol developed by Digigram has the ability to route 64 audio channels in bi-directional mode through an Ethernet network with very low latency. EtherSound™ systems can be designed using a daisy chain topology, offering buss style routing of audio channels both downstream and upstream. Such EtherSound™ systems have limited redundancy options and require pre-configured connectivity. EtherSound™ ES-100 compatible systems can also be designed using a redundant ring topology.

CobraNet™

Peak Audio's CobraNet™ (a division of Cirrus Logic) is an Ethernet compatible network protocol that offers free addressing of bundles of audio channels from any location to any destination.


The only constraint factor, as with any network technology, is the network's bandwidth. In modern systems Gigabit Ethernet technology is used, allowing several hundred high quality audio channels to be transported over the network. Within this bandwidth, audio connections can be made completely separate from the physical cabling, allowing 'no brainer' connection schemes to be used in touring systems, and a high degree of freedom to distribute I/O locations throughout a venue.

Open systems

EtherSound™ and CobraNet™ use standard Ethernet network architecture. This means suitably chosen off the shelf IT equipment can be used to build a network, taking full advantage of IT industry developments in functionality, reliability, availability and of course cost level. Both protocols are licensed to many of the world's leading professional audio manufacturers, so products from different manufacturers using the same protocol can be combined in a system without problems

Yamaha?

Yamaha adopts an open and inclusive approach, advocating the choice of a network platform appropriate to the system's requirements. The Yamaha product portfolio includes both CobraNet™ and EtherSound™ compatible products.



8 input stagerack



150m touring fiber



Amplifier rack



150m touring fiber



24 input 8 output stagerack



150m touring fiber



Digital mixer

Networked audio touring solutions

2. Three good things to know about networked audio

One: cable weight and flexibility

In conventional analog audio systems every single connection uses a copper cable. With high channel counts and cable lengths, cable weight can easily exceed 100 kilograms. With the increasing popularity of digital mixers in the pro audio industry, digital cabling such as AES/EBU is often used to replace analog cables, reducing cable weight and increasing audio quality as electromagnetic interference and cable capacitance problems are not an issue in (properly designed) digital cabling. Serial audio formats such as MADI and network protocols such as CobraNet™, EtherSound™ and OPTOCORE® have recently become popular for studio and live applications, replacing individual copper cabling with lightweight UTP (Unshielded Twisted Pair) or fiber cabling. The weight of UTP or fiber cabling is much lower compared to individual analog and digital copper cabling. Additionally fiber cabling gets rid of grounding problems.

An analog multicore cable - or a bundle of individual cables - is bulky and not very flexible. For Touring applications this means roll-out of cables requires heavy equipment, dedicated staff and limited layout possibilities.

For installations, bulky cabling requires large conduits to be installed throughout the building which is a problem especially in existing venues. In comparison, UTP and fiber cables are thin and flexible, a drum of 150 meter fiber cable weighs just a few kilograms and can be rolled up to the 'Altitude 95' restaurant on the Eiffel Tower by just one person. Installation is easy, network cables in an audio system need very little space and can be placed in an existing cable conduit.


Two: physical and functional separation

For audio networking protocols such as CobraNet™ the functional connections are separated from the physical cabling. This means that once network cabling with sufficient bandwidth has been laid out, any connection can be made without having to change the cabling. For touring this allows 'no brainer' connection schemes to be used: just connect i/o equipment to anywhere in the system and press the power button. For installations the inevitable system changes after a project's opening ceremony only require a little programming time to change the network settings, with huge savings on cabling work as a result.

Independent from UTP and fiber cabling design signals can reach even the most remote locations in a network. It no longer matters where inputs and outputs are connected to the audio system, any UTP or fiber socket will do. In a live touring situation this allows small groups of inputs and outputs to be distributed all over the stage instead of using bulky centralised connection boxes. For installations this means more freedom of choice to use multiple i/o locations in a venue, not limited by physical cabling constraints.

Three: control!


Using network information technology to distribute audio has the advantage of including.... information technology. Control signals can be included on the same UTP or fiber cabling, so there's no longer a need to lay out additional GPI, RS232, RS422 or RS485 cables. Examples are IP video connections, software control over Ethernet, machine control using RS422 serial ports etc. If you feel lucky you can even connect an Internet modem or a wireless access point...



Analog cabling (16 ch)



UTP cabling (>100 ch)



Networked Audio connectors
with bundle selection

3. Three things to take into consideration in networked audio

One: latency

The building blocks of Ethernet networks are cables and switches. To be able to route information over a network, a switch has to receive information, study the addressing bits and then send the information to the most appropriate cable in order to reach the destination. This process takes a little time, up to 120 microseconds in a 100Mb network. As networks grow larger so does the number of switches a signal has to travel through, increasing the delay with every switch. In medium sized live audio systems the network, AD/DA conversion and DSP each cause roughly 1/3rd of the total system's latency. The total system latency must be considered and managed carefully to ensure the best sound. 'In-ear monitor' applications are the most demanding and least tolerant of latency of any kind; a latency between about one millisecond and 5 milliseconds can cause unpleasant comb filter effects, and above 5 milliseconds latency can be perceived as reverb, then echo or slap at higher values. For PA FOH and monitor speaker systems the problem is relatively small, a one millisecond increase in latency corresponds with placing a speaker just 30 centimeters further away.

The EtherSound™ protocol uses a very low deterministic (variable to the network path) latency, while the CobraNet™ protocol uses a range of fixed latency values for medium, large and very large networked audio systems to allow free routing over the network.

Some manufacturers such as OPTOCORE® and Riedel use a proprietary architecture tuned to function with low latency. These are closed systems, they only work with their own equipment.

Two: redundancy

In an analog system the audio signals run through individual cables, so if a cable breaks down typically only one connection is affected. In many cases some spare connections are planned in multicore cables so system functionality is not seriously affected if something happens and a solution is easy to accomplish.

In a network however, the failure of a single long distance cable can disable the complete system, giving the engineer a hard job restoring it. This is why networked systems have to be designed with redundancy mechanisms: the system should include redundant connections that take over system functionality automatically if something goes wrong.

Some excellent redundancy features have been developed by the IT industry in the past years as banks, nuclear power plants and space agencies also need redundancy in their networked systems just as we do. Cables can be laid out double for all crucial long distance connections; if one cable fails the other takes over.

Especially in touring applications it is advisable to use redundant hardware as well, as IT equipment is primarily designed to be used in air-conditioned computer rooms, and may be more vulnerable when used in harsh on-the-road conditions.


Three: complexity

For every functional connection in an analog system the physical form of the connection is visible, normally as an XLR cable. Anyone looking at the system, or making his way through the spaghetti wiring hanging out of the back of a mixing console, can work out what is connected to what. In a network it's quite different as the functional connections are completely separated from the physical connections. Looking at a networked system a troubleshooter only sees devices connected to other devices with a few UTP or fiber cables. One cable can carry maybe two audio signals, or three hundred and sixty eight - there's no way to tell.


Where analog systems allow DIY - Do It Yourself - design and assembly by inexperienced users - Networked audio system design requires experienced system engineers who are up to date with networking technology. This drastically changes the role system integrators, system owners and system users play in the process of purchasing, designing, building, maintaining and using audio systems, a new role everybody in the process has to get used to.



Switch causing latency



Network architecture



Networked audio design

4. What is an Ethernet network?

Ethernet

Back in the seventies the Palo Alto Research Center in California, USA (www.parc.com) developed some nifty computer technology such as the mouse, the laser printer and computer networks. From the first versions of networks such as Aloha-Net and ARPA-Net the Internet has evolved. Robert Metcalfe, first working at PARC and later founding his own company 3COM, developed a practical networking standard for use in offices called Ethernet. More than 30 years later the whole world is using this standard to build information systems, and virtually all personal computers sold today have an Ethernet port built in. The Ethernet protocol is standardised as 802.3 by the IEEE standards organization.

Building blocks

The basic building blocks of Ethernet networks are network interface cards (NIC, built into devices such as computers, digital mixers), cables to connect them to the network, and switches; devices that tie all cables in a network together and take care of the correct routing of all information through the network. The operating speed of these building blocks, determining how much information a network can carry, has evolved from 10 Megabits per second in 1972 to one Gigabit per second and higher in 2006.

Addressing

Ethernet works by dividing information streams into small packets and then sending them over the network to a certain receiver address specified by the sender.

Every Network Interface Card (NIC) has an address, and switches, keep lists of addresses connected to the network in their memory so they know where to send packets. Every NIC in the world has a unique Media Access Control (MAC) address programmed by the manufacturer. There are 280 trillion different MAC addresses, and there is only one company in the world, the IEEE standards organization, that allocates these addresses to manufacturers. This way all MAC addresses of all NICs in the world are unique: there are no doubles, the system always works.

In addition to MAC addresses, a 'user definable' addressing layer is used to make network management easier for local networks. This additional user address is called the Internet Protocol address, shortnamed 'IP' address. The IP address is always 4 bytes long, divided in a network number and a host address. This division is determined by a key that is also 4 bytes long called the 'subnet mask'; every bit of the IP address that has a 1 in the subnet mask belongs to the network number, all bits with a corresponding zero belong to the host address. The trick is that only NICs with the same network number can exchange information with each other.

In most cases the network number of small office networks is 3 bytes long and the host address is one byte. One byte (8 bits) can have a value between 0 and 255. In network setting displays on personal computers the software fills in the IP and subnet values as four decimal numbers (0-255) corresponding to the four bytes in the address and subnet mask.

In small office networks the subnet mask often has the default value of 255.255.255.0 - giving the network administrator 255 host addresses to use as only the last byte can be changed and assigned to devices on the network. The first three bytes do not change and are the network number. For larger networks the subnet mask can be changed to make room for more host addresses.


Normally users have to program the NIC's IP address manually to make the network work, but in many cases a centrally located device (switch, router or computer) can be programmed to do this automatically whenever a NIC is connected using the Dynamic Host Configuration Protocol (DHCP).

VLAN


The Ethernet 802.1q standard allows for Virtual Local Area Networks (VLANs) to be created within one high speed network. This way multiple logical networks can co-exist using the same hardware, for example to create multiple audio networks for increased channel counts. Most managed switches support the VLAN standard.

Networked audio


Every Ethernet compatible networked audio device, such as CobraNet™ and EtherSound™ devices, has an NIC built in so it can send and receive information on an Ethernet network. The audio protocols use the MAC addressing layer to send and receive data. As MAC addresses are unique the devices will work with any Ethernet network worldwide.



Robert Metcalfe's first Ethernet drawing



RJ45 type NIC port



XP NIC network settings

5. Network topologies

P2P

Strictly speaking a Point to Point (P2P) topology is not a network, although a network can be used to create such a system. A P2P system includes only two locations with a fixed multichannel connection. Examples of digital audio formats for P2P systems are AES/EBU and MADI. A distribution device such as a splitter or a matrix router can be used to include more locations in the system.

Daisy chain

Daisy chain is a simple topology that connects devices serially. The EtherSound™ protocol allows connections to be made using a daisy chain topology, with devices that read and write audio channels in a bi-directional datastream at a fixed bandwidth of 64 channels in both directions. An advantage of this topology is that the routing of network information is relatively simple and therefore fast; a daisy chained EtherSound™ device adds only 1.4 microseconds latency to the network. A disadvantage of daisy chain topology is the system behavior in case of a failure of a device in the chain: if one device fails the system is cut in two parts, without any connection between the two. EtherSound™ daisy chains can be split using switches in a star topology, but in that case the audio data can flow through the system's switches in one direction only.

EtherSound™ licensees include Allen & Heath, Arcean, Audio Performance, Auvitrans, Bitner Audio, Bouyer, Camco, DiGiCo, Digigram, InnovaSon, Martin Audio, Fostex/Netcira, Klein & Hummel, Nexo, Peavy/Crest Audio, Richmond Sound Design, TESI, VTG Audio, Whirlwind, Yamaha.

Ring

A ring topology is a daisy chain where the last device is connected to the first forming a ring. As all devices connected to the ring can reach other devices in two directions, redundancy is built in: if a device fails only that device is disabled. For additional redundancy a double ring can be used. OPTOCORE® offers a proprietary system using a ring topology with a high bandwidth of up to 500 audio channels, video and serial connections. The EtherSound™ ES-100 standard supports a redundant ring topology offering 64 audio channels.

Star

As a star topology makes the most efficient use of a network's bandwidth, most information networks are designed as a star. The center of a star carrying the highest network information traffic can be designed with extra processing power and redundancy, while the far ends of a star network can do with much lower processing power.


Variations of a star topology are 'tree' and 'star of stars'. A star topology also offers easy expansion, new locations can be connected anywhere in the network. A downside is the important role of the star location as all network information to and from connected devices runs through it; if it fails a big portion of the network is affected. A network using a star topology can be made redundant using the Ethernet Spanning Tree Protocol.

CobraNet™ uses a star topology, supporting full redundancy by offering double links to the network.


The list of manufacturers offering CobraNet™ devices include Alcorn-McBride, Ashly, Biamp, BSS, CAMCO, Creative, Crest, Crown, DBX, Digigram, DigiSpider, EAW, ElectroVoice, IED, JBL, LCS, Peavy, QSC, Rane, Renkus Heinz, Symetrix, Whirlwind, Yamaha.

Selecting a topology


For every individual application one or a combination of these four topologies is most appropriate. Decision parameters include the number of locations, channel count, latency, desirable system costs, reliability, expandability, open or closed, standard Ethernet technology or proprietary systems etc. To make a decision on choosing the topology, a certain degree of expertise on networking technology is required, often found in an external consultant or a qualified system integrator with a track record in designing networked audio systems.




P2P topology (MADI)



Daisy-chain/star topology (EtherSound™)



Ring topology (OPTOCORE®)



Star topology (CobraNet™)

6. Redundancy concepts

Trunking

The Ethernet IEEE 802.1ad link aggregation standard allows managed switches to be connected with 2 or more cables, distributing information traffic over the cables. This function is also called trunking. A big plus of such a system is that if one cable fails, the other cables take over the lost connection automatically. The aggregated link will switch to a lower speed as it misses one cable, so aggregated links should be designed with ample headroom. Trunking only makes the connection redundant, if one of the switches fails the devices attached to that switch will be disconnected.

Ring

A ring is basically a daisy-chained collection of devices with the last and the first device also connected, forming a ring. Every device is then connected to the network with two cables, so if one cable in the system fails the connection is still intact. A second failure will cut the network in two. A redundant ring topology offers excellent redundancy requiring less cables compared to star topologies.

Spanning tree

In star networks packets of information are sent through the network based on the IP and MAC addresses. It's vital that the network has a logical architecture: for every source-destination combination there can be only one path through switches and cables. If there are more paths loops can occur, with the danger that information packets can flow through the loop forever, disturbing or more likely disabling the network.

So loops are normally not allowed in star networks, with the exception for networks using managed switches that support the IEEE 802.1w Spanning Tree Protocol, shortnamed STP. Switches supporting STP can block ports that cause a loop but unblock it when the active port in the loop fails. Several loops can be created in a network to protect network areas. For full redundancy a network can simply be built double, with double switches at all locations connected to each other. The advantage is that the system can recover from any failure, the downside is that this takes a while: up to 30 seconds for large networks. Recently the IEEE 802.p Rapid STP protocol has been developed offering take-over times down to 100 milliseconds. Most managed switches support some form of STP.

Meshing

Meshing is an extreme case of using STP: every device is connected to all other devices. The result is a network that is virtually immune for sabotage: you can take out many cables before a device is isolated from the network. A downside is the high amount of cables and switch capacity required, increasing complexity and costs.

CobraNet™ Dual Link

Each CobraNet™ device has two Ethernet ports built in, labelled 'primary' and 'secondary', that work in redundant mode. Normally the primary port does the job, but if that connection fails the secondary port takes over automatically.

This protects the cabling from the device to the network, but not the network. However, the dual link ports allow each CobraNet™ device to be connected to separate switches, allowing fully redundant STP configurations including redundant switches to be used.


EtherSound™ ES-100 PPM

The EtherSound™ ES-100 standard allows devices to be connected using a ring topology, appointing one device as 'Preferred Primary Master'. This PPM device blocks the ring in normal operation, and unblocks it when the ring is broken somewhere; a function similar to Spanning Tree.

Selecting a redundancy concept

For every individual application one or a combination of these redundancy concepts can be selected. One decision parameter is the required redundancy level; in touring applications it would be sensible to use redundant switches, in installed systems single switches might be enough. Normally the minimum is to have the long distance cabling redundant, with the cables separated physically as much as possible. Another decision parameter is the recovery timing - the time the system needs to recover from a cable break or switch failure.

If a closed system such as OPTOCORE® is used, the redundancy concept is selected by the manufacturer. If standard Ethernet equipment is used, some advanced knowledge is required to select the redundancy concept and program all switches in a networked audio system.



7. Cabling

UTP cables

Most Ethernet networks are built using cables containing eight copper wires twisted by pair. The shielded variety is shortnamed STP for Shielded Twisted Pair, offering protection from electromagnetic interference. The more commonly used unshielded variety is shortnamed UTP for Unshielded Twisted Pair. These cables and their connectors come in different qualities for different applications, standardised by the Telecommunications Industry Association (www.tiaonline.org) in categories 1 to 6. The categories differ by the materials used and the twisting of the wire pairs per meter. CAT3 is a low quality cable used for low speed 10Mb Ethernet networks. For 100Mb Ethernet based networks CAT5 or higher must be used. Caution: CAT3 cable looks the same as CAT5, so always study the category indication on the cable's sleeve. For use with Gigabit systems an improved version of CAT5 is available: CAT5E. The recently introduced CAT6 has even better performance characteristics. The TIA categories are backwards compatible. Within cable categories different qualities are available: massive core for installation, flexible core for patching, protection jackets and Shielded Foiled (S/FTP) for road proof touring.

UTP Connectors

Copper Ethernet network cables use RJ45 form factor connectors. The industry mostly sells cables and connectors separately, system integrators and installers can assemble cables using simple cable tools. Installation cables (solid core) and flexible cables (stranded core) require different versions of RJ45 connectors. Switch manufacturers mostly refer to a CAT5 copper network connector as 'TX', i.e. '100BASE TX'. In the audio industry Neutrik's EtherCon[®] is often used for RJ45 road proof connector systems.

Fiber cables

Optic fiber cables can handle much higher frequencies compared to UTP cabling while cable runs of over 10 kilometer can be used. The industry offers two kinds of fiber system: multimode and singlemode. Multimode fibers can handle gigabit connections of up to 2 kilometer. Single mode fibers require a more expensive laser diode, but can handle connections of up to 80 kilometers. Both varieties are commonly available in IT shops as installation fibers; some companies such as Fiberfox[®] offer military spec fiber cables for road proof touring applications.

Fiber connectors

Fiber cable connectors come in many varieties named SC, ST, LC etc. As it's very difficult to assemble fiber connectors, cables are mostly sold including connectors. Switches often use modular systems to offer fiber connectivity; industry standards for these modules are GigaBit Interface Converter (GBIC) and its mini version called Small Formfactor Pluggable (SFP). Switch manufacturers mostly refer to fiber network connections as 'FX', 'LX' or 'SX', e.g. '100BASE FX'. For road proof connectivity Neutrik developed the OpticalCon[®] connection system offering extra protection of the vulnerable fiber connectors. Connex offers Fiberfox[®], a connection system that uses lenses to disperse the fiber signal to make it less sensitive for scratches and dirt.

Media converters

A switch without a fiber module can be used to work with a fiber connection using a media converter. Media converters are widely available for both 100Mb and gigabit connections.



RJ45 connector



Neutrik EtherCon[®]



SC fiber connector



Fiberfox[®] EBC52



Media converter

8. More about CobraNet™

Bundles

CobraNet™ devices send and receive audio in small packets over a standard Ethernet network. The CobraNet™ audio packets are called ‘bundles’, each containing up to eight channels of audio and a bundle number that specifies where the packet has to be sent to. There are 65279 bundle numbers available for a CobraNet™ network. CobraNet™ devices have several methods of setting bundle numbers for send and receive: by software on a computer using a USB or Ethernet connection with the CobraNet™ device, or by dip-switches on the device.

Unicast and Multicast

Bundle numbers can be programmed to be Unicast and Multicast. Bundle numbers 1 to 255 are Multicast, which means that they will be sent to every destination in the network; the bundle can be picked up anywhere. Bundle numbers 256 to 65279 are unicast, which means that they are sent to one destination only. Some CobraNet™ devices such as the DME Satellite series and MY16-CII can send a unicast bundle to up to four destinations, this is called multi-unicast.

Quality and latency modes

Bundles contain up to eight channels of uncompressed audio with sample sizes of 16, 20 or 24 bits. The sample rate is typically 48 kHz, supporting 96 kHz as well. CobraNet™ uses fixed latency modes of 1.33 milliseconds for mid-sized systems, 2.66 milliseconds for large systems and 5.33 milliseconds for very large systems.

The fixed latency assures that the latency is virtually the same for all connections in the system, no matter how far the signal has to travel. The quality settings (sample size, sample rate) and latency mode determine the packet size; some combinations are not possible.

The conductor

The CobraNet™ device assigned to be the system’s timing generator is called the conductor. The conductor sends out a small multicast ‘beat packet’, available all over the network with an extremely short latency. All other CobraNet™ devices synchronize their wordclock generators to this beat packet so all signals in the system are synchronized to the conductor’s timing. After receiving the beat packet, all CobraNet™ devices in the system send their corresponding audio packets right away, but wait for a fixed amount of time to receive them, allowing the Ethernet network to sort out the routing without disturbing the audio timing. A Conductor Arbitration procedure takes care of the selection of the CobraNet™ device that is timing master for the network. This is an automatic process so if the master device is removed from the network a new conductor will be appointed automatically within a few milliseconds.

Bandwidth

All CobraNet™ devices have a 100Mb NIC built in, capable of sending and receiving up to four bundles adding up to a maximum of 32 channels bi-directional. The cables from a network to a CobraNet™ device can handle a maximum of 64 channels bi-directional.

The network itself can be built with gigabit Ethernet equipment with a much higher bandwidth. Unicast bundles only use bandwidth on the switches and cables the bundle travels through on its way from the sending device to the destination device, leaving the bandwidth of all other devices untouched. But when a multicast bundle is sent on the network it travels to all connected devices, using up 1/8 of the bandwidth of all devices in the network. This means a maximum of 8 multicast bundles (64 channels) can be used in a CobraNet™ network, but many more Unicast bundles. To use more than 8 multicast bundles and/or additional IP services in a system, managed Gigabit switches with multiple VLANs can be used.

CobraCAD and Discovery


CobraCAD is a software program to design CobraNet™ systems and check if a network’s bandwidth is sufficient for the amount of bundles in the design. The software allows systems to be designed with a graphical user interface, and then advises on the bundle numbers to be used. All CobraNet™ devices and recommended switches are included in the software’s component lists.

Discovery is a software program to monitor a network’s CobraNet™ devices and check for errors in the audio streams. It can also be used to configure a CobraNet™ device and to generate report files containing all of its settings.

Both software packages are freeware, downloadable from www.cobranet.info



CobraNet™ Integrated Circuit



CobraNet™ Interface diagram



CobraNet™ licensees



9. More about EtherSound™

Topology

An EtherSound™ device receives and sends audio in small, fast packets over an Ethernet network, using up the whole 100Mb bandwidth available on a 100Mb link. The protocol supports 64 channels of 24 bit 48 kHz audio in two directions; downstream and upstream, plus a small bandwidth control channel. In the downstream direction the audio channels are sent in broadcast packets. At multiple points in the daisy chain the datastream can be looped back as unicast packets so channels can be sent back 'upstream' to previous devices, creating daisy chain sections with bi-directional connectivity.

To make this work EtherSound™ systems use a daisy chain topology; each device is connected to a previous device using the 'IN' connector, and to the next device using the 'OUT' connector. The first device in the daisy chain is called the 'Primary Master, starting a 100Mb 64ch audio stream flowing down the daisy chain, and, if used in bi-directional mode, receiving the loopback audio stream flowing up the daisy chain. Switches can be used in the system to split a daisy chain into 2 or more daisy chains. In that case the audio can flow through the switch downstream only, not back. The bi-directional segments are programmed by setting 'Loop Back' and 'End Of Loop' mode in the appropriate devices.

Routing audio channels

All devices in the system read out the packets of both audio streams, take out some channels to output as audio ('slave devices'), replace channels using audio inputs ('master devices'), or both (master/slave devices).

After inserting audio channels, the broadcast packets are sent downstream to the next device; unicast packets upstream to the previous device.

Latency

Because EtherSound™ devices have only one downstream and upstream device to send packets to, addressing is ignored and packets are forwarded almost immediately after receiving them, achieving a very low latency of only 1,4 µs per EtherSound™ device. The system uses a 5 sample buffer for synchronisation, corresponding to 104 µs latency. For each 100Mb switch in the system around 22 µs latency is added, for every Gigabit switch 2,2 µs. Adding up all these values, the total latency can be calculated for each connection.

Redundancy

In a pure daisy chain topology EtherSound™ is very vulnerable to errors: any problem in a cable or device will cut the system in two parts. Using managed switches the long distance cables can be protected by using Ethernet trunking. Auvitrans offers a dedicated unit for redundant long distance cabling with very fast take-over.

The new ES-100 standard announced in 2006, allows a redundant ring topology, offering total system redundancy. The last device's OUT connector is connected to the Primary Master's IN connector creating the ring. By setting the 'Preferred Primary Master' mode in the Primary Master device this connection is blocked during normal operation, but unblocked when a connection is lost in the daisy chain, a process similar to Spanning Tree.

Bandwidth


EtherSound™ sends the audio in small broadcast packets. This means that to transmit the full 64 audio channels upstream and downstream a huge amount of packets are travelling over the network. EtherSound™ devices are designed to be able to handle this without problems, but switches used in EtherSound™ designs must also be able to handle this processing; it is advised to consult the list of tested switches on the www.ethersound.com website. Long distance links supporting more than 64 channels bi-directional and IP services can be created using managed Gigabit switches with multiple VLANs.

CobraNet™ or EtherSound™ ?

CobraNet™ and EtherSound™ are both Ethernet compatible protocols, with many suppliers of audio and networking equipment to choose from. Each protocol offers its specific advantages and constraints. Summarized on basic topics:

Topic	CobraNet™	EtherSound™	
		V2.09	ES-100
Topology	star, tree	daisy-chain	daisy chain & Ring
Redundancy	full network	long distance links only	full network (ring)
Routing	addressed	bus style	
Network latency	low (< 1.4 ms)	very low (< 0.14 ms)	


There are many other details that need to be considered for every individual case. We advise to always keep design options open to both protocols.



Daisy chain topology



Auvitrans AVY16-ES Mini YGDAI card



Auvitrans AVRed-ES redundant cabling unit



10. System engineering

System users

From the user's point of view a properly designed networked audio system is hassle-free; offering easy connectivity and flexible logistics, supporting the most complex and demanding applications such as installed systems in theaters, concert halls, leisure centers, community centers, schools, etc. Also live touring applications such as touring theater productions, pop concerts, musicals, operas etc. using their own systems or hiring systems from a rental company can benefit from networked audio systems.

System engineering

In the engineering stage normally a part of the system engineering process is handled by the system owner's technical staff, and the other part is taken care of by a consultant or a system integrator. As network engineering requires in-depth expertise on networking technology normally not found amongst audio engineers the role of qualified consultants and system integrators will increase to cover the network specification, design and programming of networked audio systems, as well as the design of easy operating and set-up procedures for the system's users.

System specification

First of all an audio system's specification must be set up.

Audio networking opens up a wide array of new possibilities, but the sky is still not the limit; without a deep understanding of networking technology it's very difficult to assess if specifications are feasible or not. A system specification includes the number of audio channels, the number of locations, the distances between the locations, the required audio quality settings, redundancy level, control services etc. If an installed system uses an existing IT infrastructure, the IT system administrator should be included in the specification process as well. For touring applications special handling specifications should be included such as cable and connector quality and connectivity standardization.

System design

Based on the system specifications a network format, networked audio format, network topology, redundancy and connectivity can be selected that suits the specifications as much as possible.

Audio components

Closed systems offer a choice of audio components set by the manufacturer. For open systems any brand of audio component that is compatible with the audio network standard can be included. Current examples of open networked audio systems are CobraNet™ and EtherSound™.

Yamaha offers a selection of CobraNet™ and EtherSound™ compatible audio components. Alternative audio components matching the Yamaha Mini-YGDAI connectivity standard are available for MADI and A-net™.

Network components


For closed systems the manufacturer supplies the network hardware. For open systems the choice of network components is overwhelming; the mature IT market offers many brands of different quality and functional levels of network equipment. We strongly recommend to run extensive network tests for every design.

Future expansion


Closed systems offer expansion using a limited choice of the manufacturer's hardware expansion options. Open systems using standard networking technology offer user-definable scalability; after the purchase of a system both network and audio components can be added, not restricted to the component brands used in the original system.

Qualified system integrator


All networked audio system designs require an amount of system engineering to be taken care of by a qualified consultant or a qualified system integrator. There are no standards to these qualifications other than having in-depth knowledge and experience in networked audio engineering and a track record of producing system designs being used in the market.



Engineering process



CAD design (StarDraw)



CAD design (CobraCAD)

11. Investing in a networked audio system

System costs

The total cost of a system is the sum of component costs and the labour costs needed to design, build and support the system. Basically a networked audio system increases component costs and decreases labour costs.

The investment in a networked audio system also has influence on the costs of usage and maintenance of a system after it has been delivered. Using networked audio systems in the touring industry will allow for significant cost savings in logistics and set-up time. Installed systems can benefit from the low cost of ad-hoc system changes.

Component costs

A network basically replaces analog long distance cabling with a digital network. This means that the component costs of all long distance cabling are replaced by the cost of the network equipment and cabling plus the required additional audio components.

Network equipment and cabling in modern times are relatively affordable. The main costs lie in the additional audio equipment, especially when a system includes digital mixing consoles that already offer analog I/O functionality. Altogether the component costs of a networked digital mixer/stage-rack solution is more or less equal to a pure analog mixer/analog multi-core solution.

To achieve a realistic cost comparison the costs of a networked audio system should be compared to an analog system of equivalent audio quality.

Labour costs

For installed systems the labour costs can be reduced significantly as the venue's long distance cabling is reduced to a few CAT5 or fiber cables. For touring systems the labour savings will occur after the system is delivered; the storage, transport and roll-out of a fiber cable is much more efficient compared to an equivalent analog multicore cabling system.


Competitive benefits

A networked audio system has a much higher quality and functionality level compared to an analog system. As projects grow to be more and more complex every year, an increasing number of jobs simply can not be done anymore without the use of networked audio systems, giving the investor in such systems a clear competitive advantage over analog solutions. This competitive advantage should also be included in cost calculations.


The bottom line

Every system has its own economics, there are too many variables to propose sensible basic rules on cost comparison. In general when replacing an analog design with a digital networked audio design the component costs will be more or less equal, the labour costs will decrease and the competitive benefits will increase. On average the total sum of cost decrease / increase plus competitive benefits will break even starting at medium sized install and touring systems. The larger or more complex the system, the higher the cost savings.

Component cost savings



No analog multicore needed



No break out, stage snake & splitter needed

Networked audio system costs



Investment in network equipment & cabling



Investment in I/O equipment

Labour cost savings



Touring: save on transport, roll-out




Install: save on cabling work

Competitive benefits



Increased functionality - flexibility



Increased audio quality

12. Networked audio glossary

AES/EBU

A digital audio format standardised by the Audio Engineering Society and European Broadcasting Union as AES3. Uses balanced copper cabling with 2 channels per connection.

Bridge

A network device used to connect networks together. Bridges work with MAC addresses, they ignore IP addressing. To connect networks on IP addressing level a router must be used.

Broadcast

The 802.3 Ethernet standard allows information to be sent to all devices on a network as broadcast packets. EtherSound™ uses this method to send audio channels on a daisy chain.

Bundle

A CobraNet™ information packet containing up to eight audio channels with 24-bit 48 kHz quality @ 1.33 ms latency.

CAT5

Category 5 cable capable of carrying 100Mb worth of network signals over a maximum distance of 100 meters.

CAT5E

Extended specification of CAT5 cable for higher frequencies.

CobraNet™

A network protocol that uses Ethernet to transport audio as well as control and monitoring data over a network. CobraNet™ is a true network protocol, separating functional connections from physical cabling using a star topology.

Daisy chain

Method of connecting devices. In case a device fails the system is cut in two.

Dual link

CobraNet™ redundancy method by connecting a device to a network with two links; if one link fails the other takes over

End of Loop device

EtherSound™ version 2.09 and up, including ES-100, allows the creation of multiple bi-directional segments in a daisy chain. In addition to the Primary Master, any devices can be set to End Of Loop mode, blocking the upstream data.

ES-100

A new version of EtherSound™ offering increased functionality. ES-100 allows a redundant ring topology to be used.

EtherCon®

An RJ45 connector combined with road proof XLR housing, manufactured by Neutrik.

Ethernet

Most commonly used network protocol in the world, standardised by the Institute of Electrical and Electronics Engineers as the IEEE802.3 standard.

EtherSound™

A network protocol that uses Ethernet to transport audio as well as control and monitoring data over a network. EtherSound™ uses a daisy chain topology with a fixed bandwidth data flow and deterministic (variable to the network topology) very low latency. A new version of EtherSound™ with increased functionality has been introduced in 2006 as ES-100.

Fiberfox®

A road-proof system to connect fiber cables, dispersing the light signal with a lens to increase the contact surface of the connector. The large surface is less sensitive to scratches and dirt.

Fiber

A medium used to transport information using light. There are single mode and multi mode fibers. Fibers can handle high bandwidth information streams and can be several kilometers long.

GBIC

Giga Bit Interface Converter, hot swappable modules to add gigabit copper or optical connectivity to a switch.

Gigabit

One billion bits (1.000.000.000 bits; Gb). A Gigabit link can carry one Gigabit per second worth of information; 10 times more data compared to 100Mb links (100 Megabits per second, a.k.a. fast Ethernet).

Global Address

An IP address allowed to connect to the Internet. Global addresses are allocated by InterNIC (www.internic.org) in order to keep every global address unique.

Hub

(Repeater hub). A simple network device that resends incoming packets to all ports without checking addresses. Repeater hubs can be used to connect network segments together forming one big network. Repeater hub technology is obsolete and should never be used in new systems.

IP address

Internet Protocol address, a user definable address to manage information streams on a network. The IP address includes a network number and a host number. It allows information to be routed on a local area network (office network) as well as a wide area network (the Internet).

Latency

(Network latency, forwarding delay). The time it takes for an information packet to travel from the sending device to a destination device.

Loop Back device

The EtherSound™ Loop Back device sends it's data not only downstream to the next device as broadcast packets, but also as upstream to the Primary Master device (or the End Of Loop device, V2.09 or higher) as Unicast packet , creating a bi-directional daisy-chain segment between the two.

A-L

MAC address

Media Access Control, an addressing system using a 48 bit (6 byte) address, allocated by the IEEE standards organization. 48 bits equals 280 trillion unique addresses, there are no doubles.

MADI

Multichannel Audio Digital Interface, standardised by the AES as AES10. Uses a single connection to transfer 64 channels of 24-bit audio

Managed switch

A switch with extra capabilities such as handling VLAN's, Trunking, Spanning Tree, Quality of Service, statistics gathering, error reporting.

Media converter

A device to convert a fiber connection to a copper RJ45 connection and back. Media converters are available for most fiber connectors and speeds.

Megabit

One million bits (1.000.000 bits; Mb). A fast Ethernet link can carry 100Mb per second worth of information. In this document a connection speed or bandwidth of 100 Megabits per second is abbreviated as '100Mb'.

Meshing

A topology used by H/P where all devices in a network are connected directly to all other devices. Such a network is virtually immune from failures other than device failures.

Multicast

The 802.3 Ethernet standard allows information to be sent to multiple devices on a network as multicast packets. This is one method CobraNet™ can use to send bundles to all other CobraNet™ devices in the network. The bundle can be picked up anywhere in the network.

Multi Unicast

Some CobraNet™ devices support sending a unicast bundle to up to 4 destinations. For sending bundles to more than 4 destinations, multicast must be used.

Multi mode fiber

Connections capable of handling large datastreams over a distance of up to 2 kilometers depending on the network standard. Multimode connections use an inexpensive laser type.

Network class

Categorisation of a network's subnet mask; determining what portion of the IP address is the network number and what portion is the host address. Class A: 1 byte (8 bits) network number, 3 bytes (24 bits) host address. Class B: 2 bytes (16 bits) network number, 2 bytes (16 bits) host address. Class C: 3 bytes (24 bits) network number, 1 byte (8 bits) host address. Small office networks mostly use class C.

OpticalCon®

Neutrik XLR connector housing for LC type fiber connectors, protecting the vulnerable fiber ends from scratches and dirt.

OPTOCORE®

A ring topology audio network standard capable of handling more than 500 channels, video and serial connections with low latency.

OSI model

A standardised model for network protocols published by the International Organization for Standardization ISO (www.iso.org). The OSI model defines seven layers, defining the physical form of electrical data (layer 1) up to the network service application that uses the network (layer 7). MAC addressing is defined in layer 2; IP addressing in layer 3.

Preferred Primary Master

EtherSound™ ES-100 devices can be used in a redundant ring topology, setting one device as Preferred Primary Master. This device then blocks the ring (so it's a daisy chain), but unblocks it when a connection is lost.

Primary Master

The first device in an EtherSound™ daisy chain is called the Primary Master, starting the 64 channel data stream sent downstream through the daisy chain. In bi-directional mode the Primary Master is the last device to receive the upstream data. A computer running ES Monitor software can be connected to the Primary Master's IN port to monitor and control all EtherSound™ devices in the network.

Private Address

IP address to be used for private networks without getting approval from InterNIC. Class A: 10.0.0.0-10.255.255.255, class B: 172.16.0.0-172.31.255.255, class C: 192.168.0.0-192.168.255.255. These are non-routable addresses and are restricted for use only within a local subnet.

QoS

Quality Of Service. An Ethernet functionality allowing switches to limit bandwidth of individual ports.

Redundancy

Designing networks with extra functionality to automatically recover from failures in the system.

Ring

A daisy chained network with both ends connected, forming a ring. Unlike a daisy chain, a ring that can transmit data in both directions has built in redundancy: in case of a failure all devices are still connected.

RJ11

Connector used for copper cabling in phone applications.

RJ45

Connector used for copper cabling in network applications (e.g. CAT5)

Router

Network device used to connect networks together. A router works with IP addresses and is capable of routing data between connected networks with different network numbers.

RS232

Serial connection standardised by the Electronics Industry Alliance (EIA) defining electrical and mechanical characteristics, supporting low bitrate P2P connections. In 1991 an upgraded standard RS232C was introduced.

RS422

Serial connection standardised by the Electronics Industry Alliance (EIA) defining electrical and mechanical characteristics.

M-R

RSTP

IEEE802.1w Rapid Spanning Tree Protocol, a faster version of the IEEE802.1d Spanning Tree Protocol.

Serial bridge

Serial connection within a CobraNet™ network allowing the use of the network to communicate with RS232 devices.

Serial server

Device to convert RS232 or RS422 into Ethernet and back, so serial signals can be used through a network.

SFP

Small Formfactor Pluggable, a mini version of GBIC's.

Single mode fiber

Connections capable of handling large datastreams over a distance of up to 80 kilometers depending on the network standard. Single mode connections use an expensive high power laser type.

SNMP

Simple Network Management Protocol, a standards based method of controlling and monitoring devices in a network.

Spanning Tree Protocol

IEEE802.1d Ethernet standard. A protocol for Ethernet switches to block loops in networks and reserve them for use if an active link fails.

Star

Most commonly used network topology. The center of the star can be designed with high processing power switches, while the ends of a star network can be designed with less processing power. 'Star of stars' or 'Tree' structures are also common variants of this topology.

STP

Short for Spanning Tree Protocol or Shielded Twisted Pair.

Subnet mask

A number that specifies what part of an IP address represents the network number and what part the host address.

SuperMAC

An audio network standard from Oxford Technologies, standardised by the AES as AES50. Transfers 48 channels of 24-bit 48 kHz audio through a CAT5 cable.

Switch

A network device that connects network components together. Switches are intelligent hubs, forwarding incoming packets only to ports connected to the packet's target address.

Topology

The way network devices are connected in a network. Basic structures are a ring, daisy chain, star, tree.

Trunking

Using two or more cables to connect switches supporting the IEEE802.3ad Link Aggregation functionality; allowing the use of two or more connections to act as a single higher capacity or redundant connection.

Unicast

The 802.3 Ethernet standard allows information to be sent to only one specific device on a network as a unicast packet using MAC addressing. CobraNet™ uses this method when only specific CobraNet™ devices in the network are to receive the transmitted audio bundles, using up bandwidth only on the links involved. See also Multicast.

UTP

Unshielded Twisted Pair. Most commonly used is category 5; CAT5.

VLAN

Virtual Local Area Network. A managed switch can separate network traffic into two or more 'virtual' networks using the same hardware..

Wi-Fi

Wireless networking standard IEEE802.11. Most used varieties are 802.11.b (11Mb/s) and 802.11.g (54Mb/s).

Useful websites

www.aes.org

www.aviom.com

www.cisco.com

www.cobranet.info

www.dlink.com

www.ethersound.com

www.hp.com

www.ieee.org

www.iso.org

www.internic.org

www.lightviper.com

www.medianumerics.com

www.optocore.com

www.parc.com

www.sonyoxford.com

www.tiaonline.org

www.yamahaproaudio.com

Audio Engineering Society, AES3,
MADI/AES10

A-net™

Cisco

CobraNet™

Dlink

EtherSound™

Hewlett Packard

Institute of Electrical and Electronics Engineers

International Organization for Standardization

ICANN Internet Corporation for Assigning

Names and Numbers

Lightviper™

RockNet™

OPTOCORE®

Palo Alto Research Center

SuperMAC/AES50

Telecommunications Industry Association

Yamaha

R-Z

Useful websites



The complete package

The complete package

Yamaha's expanded Commercial Audio portfolio facilitates a single manufacturer solution to the most complex of audio installation and touring challenges. We offer digital mixing and processing as well as multi-channel, networking amplification and a wide range of advanced output devices.

Yamaha System Solutions

Although we are proud of our line up of excellent quality products, we understand that a system solution includes more than just products: cabling, network technology, design tools, quality management tools etc. This document aims to support networked audio system design including examples of 3rd party components.

White paper 'An introduction to networked audio'

Yamaha Commercial Audio, 2006 - Ron Bakker, Hiroshi Hamamatsu, Tim Harrison, Kei Nakayama, Taku Nishikori, Tree Tordoff

A-Net™ is a trademark of Aviom inc. CobraNet™ is a trade mark of Peak Audio, a division of Cirrus Logic. EtherCon®, OpticalCon® are trademarks of Neutrik Vertrieb GmbH. EtherSound™ is a trademark of Digigram S.A. Fiberfox® is a trademark of Connex Elektrotechnische Stecksysteme GmbH. OPTOCORE® is a trademark of OPTOCORE GmbH.