2000-01-C028

Future Automotive Multimedia Subsystem Interconnect Technologies

Visteon Automotive Systems

ABSTRACT

For the past decade or so, automotive entertainment subsystem architectures have consisted of a simple Human Machine Interface (HMI), AM-FM tuner, a tape deck, an amplifier and a set of speakers. Over time, as customer demand for more entertainment features increased, automotive entertainment integrators made room for new features by allowing for the vertical integration of analog audio and adding a digital control. The new digital control came to entertainment subsystems via a low speed multiplexing scheme embedded into the entertainment subsystem components, allowing remote control of these new features. New features were typically incorporated into the entertainment subsystem by independently packaging functional modules. Examples of these modules are cellular telephone, Compact Disc Jockey (CDJ), rear-seat entertainment, Satellite Digital Audio Radio System (S-DARS) receiver, voice and navigation with its associated display and hardware. Figure 1.0 is a block diagram of typical entertainment subsystem. This paper discusses alternatives to the module-expansion of entertainment subsystem via low speed digital control and analog audio. Moreover, the discussion is expanded to cover future multimedia and infotainment subsystem interconnects technologies.

INTRODUCTION

Recently, great achievements have been reached in information, communication, entertainment, comfort, safety and security products. Moreover, new Intelligent Transport Systems (ITS) services, requiring state-of-the-art electronics, are appearing on the market to help drivers process information, make decisions, and operate vehicles more safely and effectively.

As a consequence, our cars will be equipped more and more with digital systems communicating and exchanging information. Whenever possible, the trend is to go towards a super-integration of these systems, but a number of them will always be distributed in different locations of the car. A transport bus is necessary as a backbone for all the bit-streams and commands flowing between them. Table 1.0 lists some of the most important applications that are already present or will be soon introduced into the automotive environment.

Safety & Security	Entertainment	Information and Communication
Road-side assistance Mayday	Radio (AM/FM/DAB)	Internet access
Panic call	Audio (cassette, CD player, MP3)	E-mail
Collision avoidance	S-DARS	Weather forecast Head-line News Stock quotes Traffic updates Paging
Antitheft system	Video (TV, DVD)	Tourist information
Traffic information	Games	Navigation
Tolling system		Car diagnosis
		Mobile phone

Table 1.0: Near-Future Vehicle's Features

This paper assesses the suitability of current mobile multimedia transport for the accommodation of these technological advances. In addition, this paper identifies system and functional requirements for future mobile multimedia transport as well as differences between existing networks such as Ethernet, IEEE 1394, Media Oriented System Transport (MOST).

CURRENT MOBILE ARCHITECTURE

In the beginning of the automobile era, the primary function of a vehicle was a reliable transport. Over a period of years, the desire for a basic transport has been coupled with the desire for comfort, convenience, entertainment, information, communication, safety and security. Vehicle manufacturers made room for the new features by allowing for the vertical integration of analog audio and adding a digital control. Despite the digital nature of most of the new added modules and the introduction of Digital Signal Processor (DSP) within the mobile multimedia system, the vertical integration of audio and its transport remained analog.

The Current mobile architecture with its analog transport has the following limitations:

- Analog transport complicates the subsystem interconnects, decreases reliability, adds weight and cost. Two twisted pairs are required for cabin media, one twisted pair is required for voice module, one twisted pair is required for cellular phone module, one twisted pair is required for navigation module, 3-5 wires are required for low speed multiplex scheme and synchronization. Additionally, two more twisted pairs are required for an optional media player such as CDJ. In the case of rear seat entertainment, more wires or coaxial cables are required for video. The number of wires or coaxial cables required for video applications is proportional to the number of rear seat occupants. Moreover, these wires require wide connectors with more pins at the analog power amplifier's input connector.

- The module level expansion strategy and vertical integration of analog audio resulted in a closed architecture with limited expansion path. The number of pins available at the power amplifier's input connector bounded the expansion path. In addition, the new subsystem had a short life and is not compatible with the digital trending of future entertainment features such as digital audio, digital video, Digital Audio Broadcast (DAB) and next generation of compact disc technology, Digital Versatile Disc (DVD).

- The module level expansion strategy and vertical integration of analog audio resulted in costly subsystem architecture. Often, modules added to the subsystem exhibited wasteful redundant hardware resources in order to achieve compatibility with an analog architecture. An example of this hardware wastefulness is the addition of a digital-to-analog converter to the output of CDJ or CD player to achieve compatibility with analog integration and processing. Moreover, the hardware resources available for each module are for that modules' own use and can't be shared with other subsystem's modules. This will add cost to each module and will contribute to the overall cost of the subsystem.

Presently, the module level expansion strategy and vertical integration of analog audio has reached its upper integration limit for an HMI, AM-FM tuner, voice, cellular phone, CDJ, a media player, Steering Wheel Control (SWC), a Rear Integrated Control Panel (RICP) and navigation. The implementation of such a subsystem requires seven modules and a minimum of 31 wires. A saving of one module and four wires is possible if a media, HMI, tuner, and power amplifier is packaged in one module. However, a complicated heat dissipation and overall power management strategy is required for the success of such integration, and may limit the audio performance.

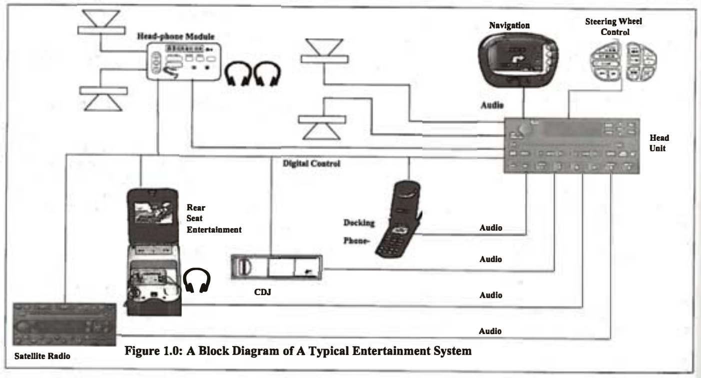

Figure 1.0: A Block Diagram of A Typical Entertainment System

REQUIREMENTS FOR NEW TRANSPORT

Vehicles are running out of the real estate required to house new modules. Interconnect harness thickness and costs are ever increasing. This section is a discussion of both system and functional requirements for a new transport:

- Open Standard: the new transport shall be an open standard to ensure that all vehicle and electronic product manufacturers have equal access to the standard and to the market.

- Minimal Standard: the new transport standard shall be extensible and capable of migration to future technologies and different physical media with no impact on application software. A layered approach to the protocol, such as is used in the reference model of Open System Interconnect (OSI) model, shall be used.

- The new transport shall be digital to allow multiplexing of data, control, audio and video over the same media. In addition, digital transport enables open system architecture; a node can be added at anytime during the vehicle's life without modifying an existing node's connector. Moreover, a digital transport enables the natural transport of digital data from one ever increasing digital application to another without exhibited wasteful redundant hardware resources such as Digital-to Analog (DAC) and Analog-to-Digital (ADC) converters.

- Safety & Security: the new transport shall provide an environment into which devices can be plugged, unplugged, and operated in a vehicle in a manner which does not threaten the integrity of the vehicle or the devices that are being used. The new transport shall include suitable security for transactions involving the exchange of money and/or proprietary information.

- Manufacturability: the new transport compatible devices and software shall be easy to design, implement and integrate. This implies minimal specifications and unambiguous requirements. Ease of manufacture helps to assure the wide adoption of the standard and contributes to lowering the cost of the manufactured items.

- Ease of Use: the addition or removal of devices shall be easy enough for a consumer to accomplish the task with little or no expert assistance required.

- Low Cost: the incremental direct material cost to implement a new transport node shall be a fraction of the cost of the application supported.

- Graceful Degradation: the new transport's physical layer and its supported bus topology shall be designed such that any fault shall not cause any damage to the cable, the vehicle or any attached devices. Functional operation under any of these conditions may cease but shall resume within the boot/discovery time after the fault is removed. The new design shall be such that no single failure, other than a fault in the physical layer as described above, or the loss of primary power, shall cause the entire system to fail.

- Hot Plug and Play: it shall be possible to attach devices to, and remove devices from, the new transport system at any time, whether power is on or off.

- Self-Identification: devices attached to the new transport shall be able to identify themselves to other system devices and shall self configure to obtain unique addresses on the new transport. No user intervention shall be required to complete this configuration other than the provision of the appropriate application software.

- Short Boot/Discovery Time: the new transport and all attached nodes shall complete self-configuration within one second after device initialization. When a new node is added to the system while it is operating, detection of the new node and reconfiguration of the system to include it shall be completed within 2 seconds.

- Peer-to-Peer Communications: the set of devices that is likely to be attached to the new transport is unpredictable and no single device is guaranteed always to be present. Therefore, a device on the new transport shall be able to communicate directly with any other device in the system and the vehicle without need for any additional device. An application that is implemented across multiple devices may choose to implement a "master controller" for that application but this shall not be a requirement for all applications.

- Automotive Physical and Electrical Specifications: the new transport's physical layer shall meet automotive environmental (temperature, vibration, shock, EMI, etc.) and electrical specifications (reverse voltage, load dump, etc.) for the environments in which the devices are installed, e.g., passenger compartment, trunk, etc.

- Extensibility: The new transport shall be extensible to accommodate evolving technologies and new applications. A layered approach to the protocol is required to guarantee minimum impact on existing designs as new physical layers and new applications are developed.

- Security and Authentication Services: The new transport protocol shall provide security and authentication services for access to vehicle functions. It is anticipated that additional services will require additional security measures to accommodate applications or transactions that require billing, authentication, confidentiality, confirmation, non-repudiation, etc.

- Bit Error Rate: the bit error rate shall be less than 1 x 10^{-9}. Applications requiring better than this shall be able to implement appropriate measures at higher layers of the protocol.

- Time-Critical Delivery of Packets- Deterministic Latency: Some applications may require that messages be transmitted within a given time period. The new transport shall be deterministic and it shall be possible to determine the maximum latency for any message in a given system configuration

- Private Message Service: Equipment manufacturers wish to be able to develop applications that span their own suite of products and provide competitive functions and features not achievable when products from different manufacturers are interconnected. The new transport protocol shall support the implementation of private messages that will allow such applications to be developed.

- Power Loading: the new transport gateway shall make power available for all devices connected to the new transport system. The total operating current drain of all devices connected to the new transport shall not exceed the capacity of the gateway unless power is routed directly to the device from another source or the device is self-powered (e.g., internal batteries).

- Internetworking: it is anticipated that wireless access to and from the Internet will be required for many devices attached to the new transport system. The new protocol shall not preclude the future implementation of internetworking services across multiple gateways or bridges between the new transport and other subnets such as Intelligent Transportation Data Bus (IDB).

- Explicit Device Addressability: it shall be possible to send a message from a device to a specific other device. It shall be possible for the sender to determine and use the desired recipient's unique address to deliver this message and all other devices shall ignore it.

- Broadcast Messages: it shall be possible for an application to generate a broadcast message to all devices connected to the new transport without having to address each one explicitly. A broadcast message may or may not require acknowledgment or confirmation of delivery. For example, an application may require confirmation (acknowledgment) that at least one receiving device capable of acting on the message has received the message.

- Consumer-Friendly Device Connection (No Special Tools Required): in most cases, it shall be possible for a consumer to install new transport nodes with common hand tools. There will be cases where professional installation may be required, but this should be the exception, not the rule.

- Wake/Sleep: any node shall have the ability to wake up the new transport system or put it to sleep. It shall be possible to wake up the nodes by sending a wake-up message, such as a pager. It shall be possible to put any node and all connected devices back to sleep with a sleep message. Absence of message traffic on the new transport for more than 30 minutes shall cause the new transport and all attached devices to go to sleep.

- Priority Sensitive Flow (Isochronous): consumer electronics devices such as video games, DVD and MP3 players, Dolby AC-3 audio components, etc., may require support for high speed isochronous data communications (i.e., data packets delivered at a guaranteed rate in a guaranteed order).

- Data Types: The new transport protocol shall allow any data type (ASCII, binary, bulk, etc.) to be transmitted in a message without need for any special escape characters or other similar artifacts added by the application.

- Fair Access to the new transport system: no single device shall be allowed to monopolize the new transport system.

- Message Priority Flagging: the new transport protocol shall provide a means to specify that the current message is a high or normal priority message. The protocol simply provides a mechanism to identify the message as a high priority message. It is up to the device manufacturer to determine whether support is provided and up to the application software to determine what action is to be taken when a high priority message is received or presented for transmission.

- Confirmation of Message Delivery, if required: a device sending a message to another device shall be able to explicitly request confirmation of error-free delivery of that message from the receiving device.

EXISTSING PROTOCOLS: ETHERNET, IEEE 1394 AND MOST

Fast Ethernet

Ethernet is a term commonly used to describe a variety of network implementations that share the same basic technology. Some early varieties of Ethernet are 10Base-2 and 10Base-5 which are also called 'thin net' and 'thick net' respectively. All nodes on such network tap into a single cable. A later version of Ethernet, 10Base-T, introduces the concept of a hub or a repeater. All nodes are connected directly to a single repeater, which simplifies cabling and provides buffering of electrical signals.

A newer version of Ethernet, called Fast Ethernet, operates at 10 times the speed of 10Base-T or 100 Mbit per second. It has the same star topology as 10Base-T and comes in a few different versions called 100Base-TX, 100Base-FX, and 100Base-T4. The difference between these versions is the physical layer which is electrical for TX and T4, and optical for FX. In this paper, Fast Ethernet will refer to the most common version 100Base-TX.

Ethernet Topology

Thick net Ethernet uses a single cable as a backbone for the network. Each node in the network taps into this cable through what is called a T connector. In an office environment, this cable could be routed through the ceiling with taps dropping into each office. This works well except it is difficult to add new users to the network and signal quality is sometimes difficult to control.

A Fast Ethernet uses a central repeater which connects directly to each node. This star topology enables the signal quality on the transmission line between a computer and the repeater to be well controlled and provides a relatively simple means to add users. The repeater has a number of ports which connect to each computer on the network. To add another computer, a wire is run from a free port on the repeater to the new computer. If there are no free ports, typically, another hub (or switch) can be connected to an uplink port to expand the network to virtually any size.

Ethernet Physical Layer

In a thick Ethernet, all computers are connected to one coaxial cable. This cable is used for sending and receiving messages. When the bus is idle, the voltage on the coax cable remains in a high impedance state at an intermediate level. This level is not a one or a zero, so that all nodes can easily determine if the bus is idle. When a node is transmitting, the voltage on the bus is pulled to high and low voltages depending on the data to be transmitted. Only one computer can send information at any one time. If multiple nodes try to transmit at the same time, a collision occurs, the data from both computers is corrupted, and both computers have to stop transmitting and try again when the bus is idle.

In Fast Ethernet, all nodes connect to a central repeater through two sets of twisted pairs. One pair is for transmitting and the other for receiving. Although each node has its own cable, the network operates exactly like thick Ethernet at a higher level. When a computer sends a message to the central repeater, the repeater sends the data exactly as received to all other computers connected to it. Again, if two computers try to send at the same time, a collision occurs, and both computers must try again later.

Ethernet Arbitration

Since all the computers on an Ethernet share the transmission media, only one computer can send information at any one time. If multiple nodes try to transmit at the same time, they must arbitrate for use of the bus. The rules that every node follows is called Carrier Sense Multiple Access with Collision Detection (CSMA/CD). Before a node sends a message on the bus, it sends a stream of one's and zero's called a 'carrier'. All other nodes on the network sense this carrier and do not attempt to send their own message until the original node completes its message.

There is a finite amount of time from when a node begins sending a carrier to when all nodes detect this carrier. During this time, other nodes may attempt to send a carrier. If this happens, a 'collision' will occur which corrupts the data on the bus. Transmitting nodes will both 'detect' this condition and stop transmitting. The rules then specify that each node must wait a random amount of time before attempting to transmit again. The probability that another collision will occur is low.

Ethernet Switches

Ethernet switches have become more popular than repeaters in recent years. An Ethernet switch is a more sophisticated device which prevents all messages from being sent to all computers connected to the switch. When a Fast Ethernet is implemented with a repeater, all computers connected to the repeater are said to be on the same collision domain. This is because a repeater operates just like the coaxial cable in thick Ethernet. If two computers try to send messages at the same time, the messages collide.

An Ethernet switch divides the network into many collision domains. Each port on a switch is a different collision domain. For example, if one computer is connected to a port on an Ethernet switch, the collision domain consists of two nodes; the computer and the switch. If a port on a repeater is connected to a port on a switch, the collision domain consists of the switch and all computers connected to the repeater.

The switch has intelligence which learns the addresses of the computers connected to each port. When a message is received at one port, the destination address is determined and the message is sent out the appropriate port. Switches are typically much more efficient than repeaters. However, they cost more.

Ethernet Communication Mechanism

Information in Ethernet networks is communicated in packets. Each packet consists of a header, usable data and a checksum. The header contains information such as source and destination address, the length of usable data and possibly information about the message type. The checksum is a code sent at the end of the message so that the receiving node can determine if the packet was corrupted during transmission.

Since the header and checksum are only used to send the packet safely from the transmitting node to the receiving node, it is considered network overhead. It is not information usable by the application. In Fast Ethernet, this consumes 18 bytes. If you include the arbitration time, the total overhead is 38 bytes. In addition, the minimum usable data is 46 bytes per packet. Even if you only wish to send one byte, you still must send the 18-byte header, 46 bytes of data, and wait 20 bytes worth of time for the bus.

The efficiency of the network can be defined as the number of user data bytes per packet divided by the number of bytes in the packet plus the overhead of waiting for the bus. If only one byte of user data is sent per packet the efficiency is 1/(64+20) X 100% = 1.2%. Since the maximum user data per packet is 1500 bytes in 100BaseT, the theoretic maximum efficiency is 1500/(1518+20) X 100% = 97.5%.

The maximum efficiency is never achieved since many collisions will occur if there is a lot of activity on the bus. The effect on efficiency is difficult to predict.

Ethernet Audio Example

Let's consider a PC with a sound card connected over the network to a server which stores WAV files. A particular Wav file from the server can be played on the sound card in the following way. The client software on the PC manages a FIFO (First-In First-Out) which continually outputs audio data to the sound card. When the FIFO gets close to being empty, the client sends a message to the server to send more data. The server sends another packet to the client to fill the FIFO up again. As long as the server sends the new packet before the FIFO empties, audio can be heard. If the server responds slowly or the network is very busy, the packet may not arrive in time, the FIFO will empty and sounds will momentarily stop. This is unacceptable for most audio applications.

There is a trade off between FIFO size, the frequency of requests for more data and packet size. Since the large packets are more efficient than small packets (% overhead from header, etc), let's assume we will use the largest packet size; 1500 bytes of user data. If the audio sample rate is 48 kHz, and the audio is 16 bits/sample stereo, then we need an average of 192K bytes/second or 128 packets/sec. The overhead for the header, checksum and the required idle between packets is 38 bytes. Since the client software on the PC with the sound card must inform the server when the FIFO is nearly empty, there are another 84 bytes of overhead to send this message to the server.

The minimum total bandwidth required for one audio channel is:

1500 + 38 + 84 = 1622 bytes/packet

1622 bytes/packet * 128 packets/sec = 207616 bytes/sec

207616 bytes/sec * 8 Bits/bytes = 1.66 Mbit/sec

Since the packet size in this example is the largest allowed by Fast Ethernet, the network overhead is small compared to the audio data throughput. The disadvantage of the large packet size is the buffer size requirement in the Client. It must be 1500 bytes deep plus more for handshaking. If the extra depth is not large enough for the network to guarantee another packet will arrive prior to the buffer emptying, a loss of audio quality may occur. If the buffer empties, the audio stops. In a Fast Ethernet, it is impossible to guarantee any bandwidth. If the network has lots of traffic, you may not even be able to get the 1.66 Mbit/sec average throughput that is required. More commonly, at times of high traffic the buffer may empty no matter how large it may be.

IEEE 1394

The IEEE 1394 specification is a hardware specification that defines the serial bus architecture that Apple computer initially named FireWire. It defines serial bus specific extensions to the Control and Status Register

(CSR) Architecture for Microcomputer Buses formally adopted as ISO/IEC 13213 (ANSI/IEEE 1212).

This architecture defines a set of core features such as node architecture, address space, common transaction types, Control and Status Registers (CSR), configuration ROM format and content, message broadcast mechanism to all nodes and interrupt broadcast to all nodes. IEEE 1394 specifies how units attached to a serial bus can talk to each other, but does not define the protocols used to communicate between the nodes.

IEEE 1394 is similar to Fast Ethernet in many ways. Data is always communicated between nodes in packets. If multiple nodes try to send packets at the same time, they must arbitrate for the bus. The information in the packet headers, the packet sizes and the arbitration method, are different. However, the fundamental mechanisms are similar.

The most significant feature that IEEE 1394 provides (which Fast Ethernet does not) is guaranteed bandwidth for real time applications. These applications are allocated isochronous bandwidth which enable real time data to be communicated in packets sent at regular time intervals. This is an improvement over Fast Ethernet. However, it will be shown that there are still serious limitations.

The raw bit rate for IEEE 1394 is defined to be selectable between approximately 100, 200, and 400 Mbit/sec. Work is currently being done on an 800 Mbit specification as well. Silicon is currently being advertised to run at both 100 and 200 Mbit. However, most implementations are now at 100 Mbit/sec which is the same data rate as the large installed base of 100BaseT. Gigabit Ethernet is currently under development as well.

IEEE 1394 Physical Layer

The physical layer for IEEE 1394 consists of two sets of twisted pair wire for signals and two wires for power and ground connected between each pair of nodes. One set of twisted pairs is called data and the other is called strobe. When one node begins to send a packet, it sends Nonreturn-to-Zero (NRZ) data on the data line and transitions the strobe line only between consecutive 1's or 0's. Both sets of twisted pairs are bi-directional. Each node sends and receives data on the same sets of wires. When neither node is sending data, the twisted pairs are held in a high impedance state.

In contrast, the physical layer for Fast Ethernet consists of two sets of twisted pairs; one pair is used for transmitting data and the other pair for receiving data. This approach makes recovering data from the transmission line slightly more difficult than with IEEE 1394 since there is no strobe signal. An advantage of keeping the transmission lines unidirectional, however, is that the transmission line can be longer. The maximum cable length for Fast Ethernet is 100meters, while it is only 4.5 meters for IEEE 1394.

Each node on an IEEE 1394 bus (or in a Fast Ethernet) has its own timing source which is typically a crystal oscillator. This timing source is used by a node to transmit data and is used by a node to over sample the received data and strobe lines to recover the data. This means that all IEEE 1394 nodes are asynchronous at the lowest level. The accuracy of the timing reference in IEEE 1394 is specified to be ± 100 PPM which is typically the frequency tolerance of widely available crystal oscillators.

The nominal data rate in 100 Mbit IEEE 1394 is 98.304 Mbit. If the crystal oscillator at a particular node is operating at the high end of its frequency tolerance, it will be able to transmit data at 98.304 Mbit +100 PPM or 98.314 Mbit/sec. If it is operating at the low end of the frequency tolerance, it will transmit data at 98.294 Mbit/sec. This may seem like a trivial issue. However, the section on system timing will illustrate some important consequences for real time applications.

IEEE 1394 Topology

The physical topology for a typical IEEE 1394 network is a tree structure. Typically, a node will have a least two ports which enables multiple nodes to be daisy chained together. If a node has more than two ports, multiple branches can be created. During initialization, one node is defined to be the root node with all nodes extending down different branches. The topology can have any number of branches if no loops are created.

The tree topology, with the ability to daisy chain nodes, has the advantage of simplicity for small networks. If you have a few devices, it is easy to plug them together in a daisy chain. However, large networks can be cumbersome, particularly if network performance is optimized. To improve performance, it is desirable to minimize the propagation delay of data between any two nodes in the network. Long daisy chains can be split into many branches to reduce the delay; however, care should be taken to balance the length of the branches.

In a Fast Ethernet, a repeater or switch is required for a network with more than two nodes. This makes small networks complicated. However, it makes large networks simpler.

IEEE 1394 Arbitration

As described previously, all computers on a Fast Ethernet using a repeater (not a switch) have the same collision domain. If multiple nodes attempt to transmit at the same time, the transmitted data is corrupted, this condition is detected, and the nodes begin to arbitrate for the bus. Likewise, all nodes on an IEEE 1394 network have the same collision domain. Only one node can send a message at one time. If multiple nodes try to send messages at the same time, only one node will gain control of the bus.

In a Fast Ethernet, if a collision between two nodes occurs, both nodes will stop transmitting and wait a variable amount of time before another attempt. If they both happen to wait the same amount of time, another collision will occur. Mechanisms are built into the network to minimize the probability of nodes colliding more than once or twice.

An IEEE 1394 network operates differently. Nodes which are closer to the root node have a higher natural priority. When two nodes attempt to transmit at the same time, the node with the higher priority wins arbitration and control of the bus. In order to prevent higher priority nodes from monopolizing the bus a fairness interval is defined. During a fairness interval, all nodes are given the opportunity to send one message.

Unlike Fast Ethernet the arbitration mechanism in IEEE 1394 is deterministic. There is zero probability that nodes will collide many times before successfully sending a message. This is important for the delivery of real time data since any unpredicted delay, no matter how unlikely, may cause buffers to overflow or underflow.

The arbitration process in IEEE 1394 consists of bus request/grant handshaking between child and parent nodes. A parent node is defined as the node on a 1394 cable which is closer to the root. The node which is further from the root, is called the child. If a child and a parent both request the bus at the same time, the parent will block the child's request and send its request to its parent. The request continues down the tree until it reaches the root node. The root node issues a bus grant which travels back up the tree to the node requesting control of the bus. Once a node receives a bus grant, it can begin sending a message.

The time that it takes to arbitrate for the bus depends on the size of the network. A bus request from a node at the end of a number of branches must propagate down the tree to the root and back up the tree to the requesting node. Information must be sent down all other branches to prevent any other node from driving the bus (through a bus request), until the granted message is sent. The total arbitration delay for a network with N hops (from parent to child or child to parent) between the furthest two nodes, is about N x 80ns in a 100 Mbit IEEE 1394 network.

IEEE 1394 Communication Mechanisms

While Ethernet treats all packet data the same, IEEE 1394 provides different types of packets. The primary packet types are asynchronous and isochronous. Asynchronous packets are functionally equivalent to Ethernet packets. Isochronous packets are only available in IEEE 1394, and provide guaranteed bandwidth to time critical applications.

IEEE 1394 Asynchronous Packets

An asynchronous packet consists of a header, a header checksum, user data and a user data checksum. The header contains information such as source address, destination address, message length, message type, etc. The size of the header and the checksums is typically 24 bytes long. The user data can be up to 512 bytes long in a 1394 network operating at 100 Mbit/sec.

When a particular node receives a message, an acknowledgement signal is automatically sent back to the sending node. If the sending node does not receive an acknowledgement signal within a specified time, it will try to re-send it a specified number of times.

The total time required for sending a minimum size message consists of the arbitration time, the time to send the packet, the time to wait for an acknowledgement, the time for the acknowledgement message and a sub-action gap time. The sub-action gap time is the time required by a node to wait after the bus is idle before it can issue a bus request. The sub-action gap time is required to ensure that no node will issue a bus request before the previous acknowledgement is received.

The total time for a one byte message in a large network (16 hops) running at 100 Mbit is about 1.4 µsec for arbitration, 2.2 µsec to send the message (28 bytes * 8 bits/byte * 10ns), and 5.6 µsec for acknowledge and sub-action gap. The total time is about 9.2 µsec. This means the efficiency is about 80 ns / 9.2 µsec X 100% or 0.9%. The total time to send a maximum size message (512 bytes) is about 50 µ sec with an efficiency of about 85%.

IEEE 1394 Isochronous Packets

The most significant improvement of IEEE 1394 over Ethernet for real time audio, video and voice applications is 1394's isochronous channels which provide guaranteed bandwidth to applications. When an application is allocated an isochronous channel, it is guaranteed to be able to send isochronous packets at regular intervals to any other nodes in the network. Each node in the network has a counter which defines these regular intervals. The Isochronous Resource Manager sends a message periodically to synchronize the counters in all nodes.

The Isochronous Resource Manager has a local counter which is clocked by a local 24.576 MHz clock and counts up to 3072. When the counter reaches 3072, the Isochronous Resource Manager resets its counter and broadcasts a message to all the nodes in the network to reset their local counters. The frequency of this synchronization message is 24.576 MHz / 3072 or 8 kHz.

After the synchronization message, all nodes that have been allocated isochronous bandwidth are allowed to send their packets. After all isochronous packets have been sent, then nodes that need to send normal asynchronous packets are allowed to do so. The time between the 8 kHz sync messages is 125 µsec. The maximum time that can be allocated to isochronous data is 100 µsec. This ensures that some bandwidth will always be available for some asynchronous packets.

IEEE 1394 Audio Example

The concept of an isochronous channel resolves the problem illustrated in the previous example of communicating audio over Fast Ethernet. Since a node that is allocated isochronous bandwidth is guaranteed to be able to send a packet every 125 µsec, the FIFO described in the Fast Ethernet example will not empty due to congestion on the network. In addition, since the audio data packets are guaranteed to arrive at fixed rate, it is intended that handshaking messages are not necessary. These are required in the Fast Ethernet example since the server needs to know when the FIFO in the PC with the sound card is nearly empty.

To send 16 bit stereo audio data at a 48 kHz-sample rate over an isochronous channel requires 24 audio data bytes per isochronous packet. This is shown below:

48 kHz * 2 bytes/sample * 2 channels = 192K bytes/sec

192k bytes/sec / 8k packets/sec = 24 bytes/packet

The header for an isochronous packet consists of eight bytes which specify the length of the data (in our case 24 bytes), the channel number, etc, It also provides a checksum for just the header. The checksum for the data consumes another 4 bytes. In total, there are 12 bytes of header and checksum overhead to send the 24 audio bytes.

Additional overhead comes in the form of arbitration time, gap time between packets and packet start and end symbols. The gap time and start and end symbols consume a fixed time equivalent to the time it takes to transmit about 3 bytes in 100 Mbit IEEE 1394. The arbitration time depends on the size of the network. For a large network that has 16 nodes between the furthest nodes, the arbitration time is approximately equivalent to 20 bytes of data.

In summary, to send the 24 bytes of audio data, a 1394 network consumes up to 35 bytes of overhead. The total time required sending such as message is about 4.7 µsec and the efficiency is about 40%. Since the maximum amount of time that can be allocated to isochronous is 100 µsec/Frame, the maximum number of audio channels that a 100 Mbit IEEE 1394 network can support is about 20.

Voice is typically sampled at 8 kHz with 8-bit resolution. The minimum data field in an isochronous packet is 4 bytes; so the packet size will be at least 16 bytes. Including the overhead described above, the time required to send this message is 4.1 µsec and the efficiency is about 2%. The total number of voice channels that a 100 Mbit IEEE 1394 network can support is about 24.

If 100 µsec per frame are consumed by isochronous packets, 25 µsec will be left for asynchronous messages. As shown in the previous section, the time to send a minimum size packet is 9.2 µsec. Consequently, a maximum of two asynchronous packets can be sent per isochronous frame. In other words, 16000 relatively small asynchronous packets can be sent per second.

As shown earlier, a 512-byte packet (largest possible in a 100 Mbit IEEE 1394 network) takes 50 µsec to send. Since only 25 µsec are available per frame for asynchronous messages, it takes two frames to send such a message. The message rate will be 4000 Bytes/second. It is possible to send such message, since the isochronous packets on the second frame are delayed until after the large asynchronous packet has been sent. If the transmit time of a packet were allowed to be larger than 50 µsec or if less than 25 µsec per frame were reserved for asynchronous packets, delivery of all the isochronous packets could not be guaranteed.

IEEE 1394 System Timing

The timing source for each node in an IEEE 1394 network is typically a crystal oscillator. As described in the physical layer section, the frequency of crystal oscillators can vary by ± 100 PPM from their nominal value. At every node a local 24.576 MHz clock is generated to clock the modulo 3072 counters used to create the 8 kHz isochronous frames. The Isochronous Resource Manager must send a synchronization message every 125 µsec to resynchronize the counters in all nodes since they are all clocked by slightly different frequencies.

If the Isochronous Resource Manager has a crystal operating slightly faster than another node, the modulo 3072 counter in the other node will sometimes count to

DOCKET ALARM

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts

Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research

With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips

Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.

WHAT WILL YOU BUILD? | **sales@docketalarm.com** | **1-866-77-FASTCASE**

fastcase®
Smarter legal research.