Wheatstone Corporation

Technical Document

IP Audio Networks
Planning for Reliability

In planning a LAN (local area network) structure for use in an IP audio network, it is important to understand the physical and electrical requirements needed to assure reliable network performance. While a $30 Ethernet switch and a few cat5 patch cables from your local office supply store will suffice to stream a couple of channels of audio around, any real world studio installation needs to be carefully designed and constructed.

Audio over IP technology uses standard protocols and Ethernet switch hardware in a way they were never designed to be used. It is a testament to the flexibility of the original designs that it even works at all. Here’s why.

Ordinary computer LAN traffic is “bursty” by nature. There are long periods of relatively little traffic punctuated with short bursts of data. This data in general has no absolute timing requirement; we are simply moving data from one place to another and as long as it gets there fast enough we’re satisfied. LAN traffic consists of 3 main classes of data:

- **UDP** (User Datagram Protocol) supports broadcast traffic. One transmitter can send data to every member on the network. This is a very efficient mechanism for transmitting information that “everyone” needs to get because the information is sent once, and there is no inherent handshaking to assure that everyone receives the information. An example of a UDP message would be a computer being plugged into a network. It needs to announce to every other device on the network (without knowing what other devices may be there) that it is present and available.

- **TCP** (Transmission Control Protocol) is a unicast protocol. One transmitter sends data to one specific receiver, which in turn must acknowledge back to the transmitter that it received the data correctly. This is a robust protocol for transmitting information to a specific recipient. A example of a TCP message would be a print job being sent from a PC to a printer. None of the other devices on the network need know anything about the print data but it is very important that every bit of it gets to the printer correctly. Because at the hardware level packet collisions can and do occur, TCP provides a method for missing or corrupt data to be resent until the full message has been received correctly. TCP messages form the majority of common office LAN traffic.

- **Multicast** is a third method for information transfer. It is designed for the case where the data is of interest to more than one recipient, but not everyone. Multicast is an improvement on UDP; no handshaking or resending is involved, and the sender transmits the data only once. The key is that filtering is applied as to where the data is sent. This is the job of a managed Ethernet switch. The switch maintains a list (called a Multicast

> Table) of the recipients (called subscribers) of each data channel and forwards the information from the transmitter to only those receivers who have subscribed to it. The Ethernet switch maintains this table; adding new subscribers to a data channel as they request and removing old subscribers after they are finished with the data channel.

Multicast is the mechanism used for IP audio networking because it is the most practical. If UDP broadcasts were used, every device on the network would be flooded with every audio stream simultaneously when only one or even none were desired. Conversely, if TCP were used, the transmitting device would have to send (and dutifully wait for acknowledgment of) duplicate streams of data, one to each recipient, again flooding the network with multiple copies of the same information. Multicast allows the intelligence built into the Ethernet switch to send the desired audio streams to the waiting recipients with the least amount of network traffic.

Why is minimizing network traffic so important in audio over IP systems? The simple answer is that audio streams represent a relatively huge amount of continuous (as opposed to bursty) and timing sensitive data. In fact most people are astounded when they realize how much information is actually transmitted in an audio stream. A few examples may help:

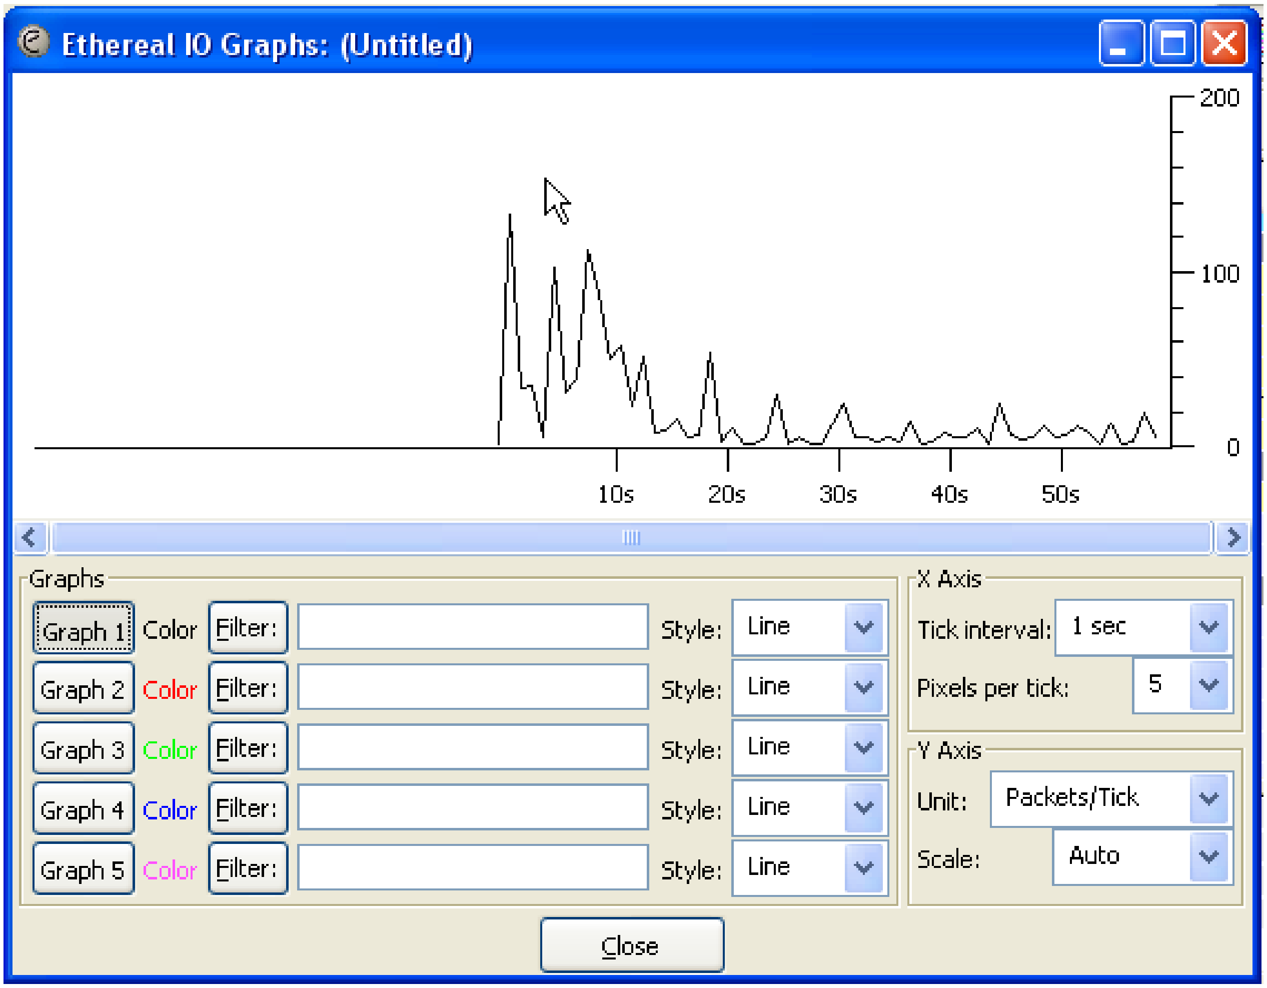

Style: Line
Style: Line
Style: Line
Style: Line
Style: Line

Tick interval: 1 sec
Pixels per tick: 5

Unit: Packets/Tick
Scale: Auto

Figure 1: Office Network

Figure 1 shows a one minute capture of all of the network traffic on a segment of an office network. This segment consists of 18 computers of various types and three printers. You can see the variable nature of the network traffic, but note in particular its peak rate of 140 packets per second; each packet represents about 1500 Bytes of data. A total of 2188 packets passed through this network segment in one minute.

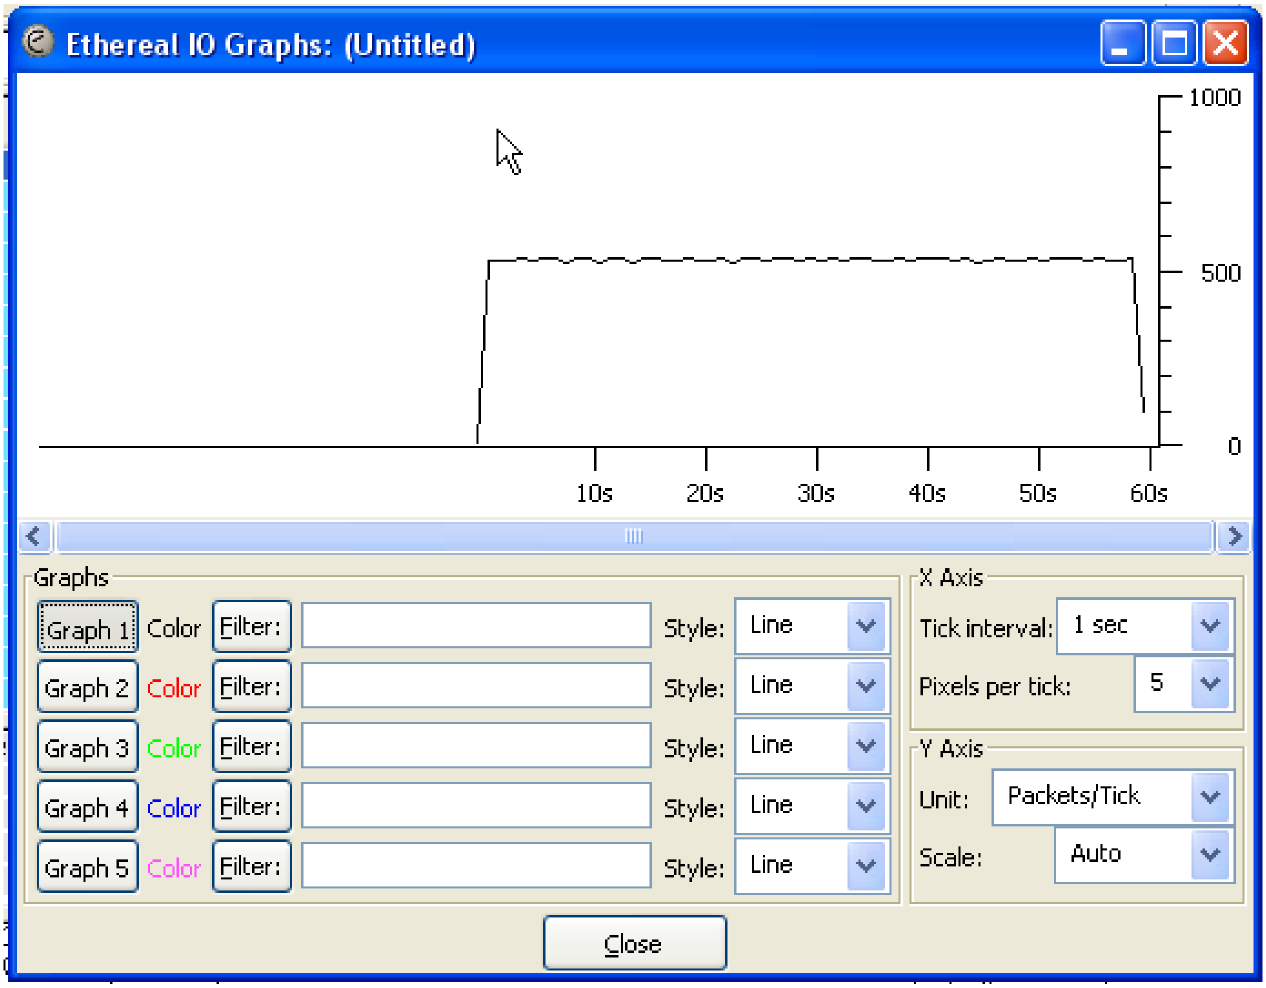

Style: Line

Tick interval: 1 sec

Pixels per tick: 5

Unit: Packets/Tick

Scale: Auto

Figure 2: AoIP Without Audio

Figure 2 shows a one minute capture of the network traffic of one port of the Ethernet switch in an AoIP system. The port being monitored was connected to a PC and no actual audio streams were going through the port, but only background and announce messages. Note that the rate of the traffic is absolutely consistent and high. A total of 32,892 packets passed through this one port in one minute.

Now look at figure 3. This again shows a one minute capture of one port, but this time the port is connected to an AoIP device that is receiving 3 stereo channels of audio. Again the traffic rate is consistent and has risen to a rate of over 25,000 packets per second. A total of 1,581,292 packets passed through this one port in one minute, or 218 MB total data, and this is for just 3 channels of audio. Clearly Ethernet switches have to work much harder than usual in an audio over IP network.

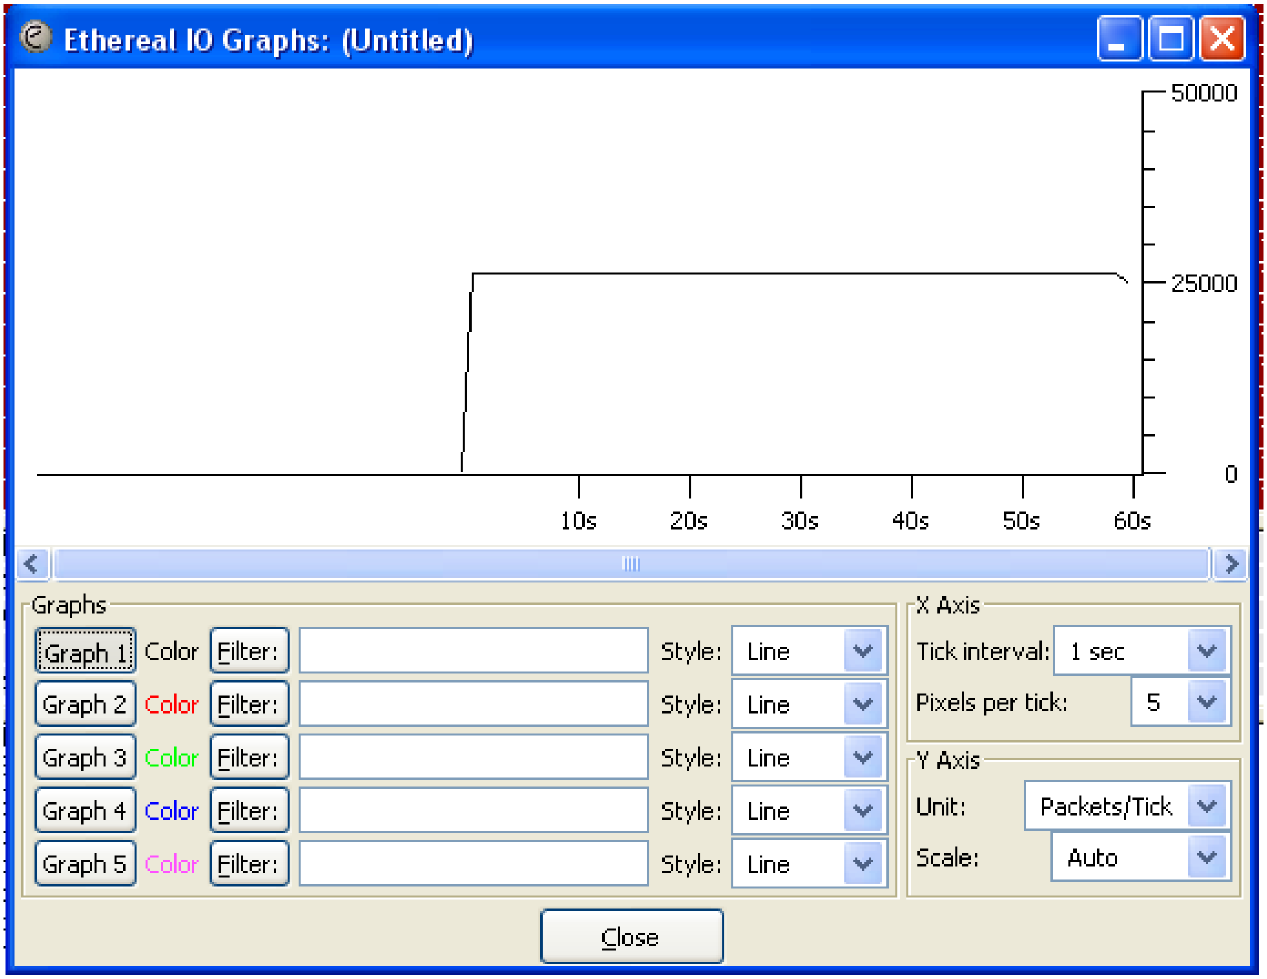

Ethereal IO Graphs: (Untitled)

Graphs

Graph 1 Color Filter: Style: Line
Graph 2 Color Filter: Style: Line
Graph 3 Color Filter: Style: Line
Graph 4 Color Filter: Style: Line
Graph 5 Color Filter: Style: Line

X Axis

Tick interval: 1 sec
Pixels per tick: 5

Y Axis

Unit: Packets/Tick
Scale: Auto

Figure 3: AoIP with Audio Streams

The point is that in these audio networks a consistent and very high rate of data transfer is going on continuously, and anything that causes a disruption in this data flow immediately causes a disruption in the audio.

Transfer disruption generally comes in one of two types; local or global. The type of disruption gives us a good clue as to the nature of the problem.

Local disruption occurs when the data flow to a particular device or group of devices on a network segment is interrupted. The interruption can be partial or complete. Complete interruption is the easiest case to analyze; the audio simply goes away. Usually it is the effect of a broken or missing cable connection, or a defective or improperly configured switch port. If the audio went away after someone was “working on the Ethernet switch settings,” suspect the configuration. Otherwise, look to the cables and connections.

Partial interruption is a little more complicated; the audio seems to come and go, or periodically break up. While intermittent cable connections or low quality cables can cause this problem, so can an overloaded Ethernet switch (one that is being sent more data than it can manage). Poor quality cables can be masked in an ordinary office LAN because the TCP protocol will keep resending data until it ultimately gets through.; once you start to stream audio over these cables, it becomes immediately apparent that something is wrong because the audio is clicking. popping, or dropping out (multicasting has no mechanism to replace or resend missing data).

A more subtle cause is the presence on the network segment of some computer or device that can't handle all of the data (refer to figure 3) being sent. One characteristic of a TCP connection (common to a PC) is that the Ethernet switch will automatically slow down the network speed if it is not getting reliable message acknowledgments from a connected device. This happens frequently with inexpensive printers, which is one reason why Wheatstone strongly recommends isolating AoIP networks from your common office LAN. Any slowdown in the network traffic manifests itself as audio breakup or dropouts. Switch port settings, because they affect network transfer speed, can also cause symptoms. This problem can manifest as sluggish audio response when faders are moved or buttons are pressed as a device tries to keep up with its TCP messages and responses while being flooded with millions of unneeded multicast audio packets.

Global disruption occurs when the entire data flow throughout the network is interrupted. Such disruption, when not due to an obvious cause such as power failure, is generally caused by network settings and/or changes. Remember the discussion about how multicasting works? Since audio networking depends on the Ethernet switch's multicast table to direct the flow of audio data, anything that affects that table will affect the audio. The bad news is that there are a number of things that can affect the multicast table.

Since the table reflects all of the devices on the network and their status, any time this gets changed, the table needs to be updated. Switches sometimes do this by throwing out the old table and building a new one. Until the new table is rebuilt, the audio goes away.

Network status changes can be deliberate, as when new devices are added, or unintentional, as when a device that was off is suddenly powered on. The ports on the Ethernet switch can be hard coded to guard against this type of disruption, but this is indirect protection at best since there is no way to lock the multicast table. This is another reason why Wheatstone recommends isolating your AoIP network. All it takes is for the wrong device to get plugged into a port to kill all of the audio. If all of the audio in the system disappears for a minute or two and then comes back all by itself, suspect the network configuration and multicast table and find out what caused them to change.

Now that we have some understanding of how IP audio networks work and the pitfalls in managing them, what's the best way to accomplish an installation and ensure its reliable operation? In a word, planning:

- Do plan on isolating the audio network from other networks and devices.
- Do use a managed switch for any system larger than one studio.
- Do buy a large enough switch(es) for the size of your system, these systems always grow larger over time.

- Do plan on using a separate Ethernet switch(es) dedicated to the audio network
- Do configure the switches properly as documented.
- Do use high quality cat5E or better cables throughout, including patch cords.
- Do have your cable system tested and certified, including patch cords.
- Do plan on performing any system maintenance or modifications at non-critical times.
- Do plan on using an experienced IT professional (one with experience in IP udio networks) for large and complex installations.

IP audio networks are an exciting technology. Small systems are deceptively easy to configure and install, but with proper planning and procedures even the largest can be a robust and reliable audio infrastructure.