help desk

sACN in large systems, Pt. 1 | By Wayne Howell

After delving into LLRP (Low Level Recovery Protocol) in LSi October, this month's column is dedicated to problems that can be encountered when using sACN in large channel count systems. But before looking at these, let's remind ourselves how sACN works . . .

WHAT IS SACN?

sACN (streaming ACN) is a lightweight streaming protocol that's both simple and easy to implement. It was developed as a cut-down version of ACN in response to the latter's lacklustre support, and is an ANSI standard (E1.31). sACN only handles level control data (DMX) - there is no support for RDM.

The sACN approach to sending DMX is very similar to Art-Net: the raw DMX is encapsulated in an Ethernet wrapper, but the data is sent using multicast. There are three ways that Ethernet packets can be sent.

'**Unicast**' is the simplest form of transmission and is the predominant form of data transfer on the majority of networks. In unicast, a single sender transmits data to a single receiver; the method is also extremely bandwidth-efficient because network infrastructure equipment, such as switches, can ensure the packet is only transmitted over the specific parts of the network to which the intended recipient is connected. Unicast is less efficient if receivers need to see identical data.

In '**broadcast**', the senders transmit data to all receivers. This is very beneficial for network management packets such as ARP (Address Resolution Protocol), where all devices must see the data. However, broadcasting can be very inefficient when a large percentage of devices don't need the data - Ethernet switches pass broadcast data to all their ports, which can saturate the network. This is why many switches have filters that can stop broadcast packets - I'll come on to that later.

In '**multicast**', the senders transmit to a group of receivers. Receivers 'subscribe', using a protocol called IGMP (Internet Group Management Protocol), to become a member of the multicast group. Each multicast group has a specific IP address so that all members can receive data directed to that group. It uses Class D IP addresses (from 224.0.0.0 to 239.255.255.255). Multicast can be more efficient than unicast when many receivers need to see the same data.

HOW DOES SACN USE MULTICAST?

Any sACN receiver that wishes to receive a particular universe will subscribe to a specific multicast address. The multicast address is 239.255.xxx.yyy, where xxx and yyy represent the sACN universe number. So, the IP address of the first sACN universe is 239.255.0.1.

The lighting console sends sACN packets to the relevant multicast address for each universe. So, imagine a truss with four moving lights, all patched to universe 1 - the four moving lights will subscribe to multicast address 239.255.0.1. The lighting console then only needs to send a single packet to that multicast address in order to control all four moving lamps.

On the face of it, that sounds great, but there are three key problems with the way that sACN uses multicast.

1. Transfer of multicast data on a network is not plug-and-play. You need to configure the network correctly - if you don't, all your efficient multicast data ends up being broadcast to all ports, defeating the purpose and potentially overloading the network. It's aslo quite hard to test if a network is correctly configured.

2. Network switches have limited resources - and multicast groups use those resources quickly. Multicast wasn't designed to be used with tens of thousands of multicast groups active on the same network. Many network switches have (an often unpublished) limit to the number of groups they can handle - on a low-cost switch, that could be 64. For example, even when you configure your network perfectly, if you transmit more than 64 universes, they will be broadcast.

3. Many network switches treat large amounts of multicast data as a denial-of-service attack and will start filtering out the data. In order to understand how to configure a network for sACN and check if it's working, we need to look how packets move around a network . . .

THE LAYER CAKE

Networking uses a model called the OSI Layer Model, which has seven layers hierarchically used to describe how data flows:

Layer 7 The Application Layer	The end users.
Layer 6 The Presentation Layer	Handles encryption and data representation.
Layer 5 The Session Layer	Management of the communication between devices.
Layer 4 The Transport Layer	Reassembly of lower level packets into correct order.
Layer 3 The Network Layer	Management of packets based on IP address.
Layer 2 The Data Link Layer	Management of packets based on MAC address.
Layer 1 The Physical Layer	The electronics, cables and fibre.

Data transmission starts at Layer 7 with the end user and flows down through the layers until it reaches the cable. At the other end of the cable, the data flows back up the layers until it reaches Layer 7. For the purposes of looking at how sACN moves, I have redrawn the table to compress the top three layers and show what each layer means to the sACN packet.

Layer 7 – 5	sACN users.
Layer 4 The Transport Layer	sACN is assembled into UDP (User Datagram Protocol) packets.
Layer 3 The Network Layer	Multicast addressing is handled using IP addresses. (IGMP).
Layer 2 The Data Link Layer	Ethernet switch manages packets based on MAC address.
Layer 1 The Physical Layer	The electronics, cables and fibre.

Up in Layers 7-5, we have sACN data that is meaningful to the end user. In Layer 4 it's assembled into packets called UDP and is ready to be sent. In Layer 3 our multicast groups are managed. That is where the clever part of multicast addressing happens. In Layer 2, our network switches decide which packets should be sent to which ports. Finally, in Layer 1 our data makes it onto the cable.

THE LAYER 2 & 3 CONUNDRUM

Many common Ethernet switches are of a type that manage packets at Layer 2 by learning the MAC addresses of the devices connected to each port, thereby ensuring the right packets are sent to the right ports. This type of Ethernet switch is (unsurprisingly) known as a Layer 2 switch or L2 switch. However, multicast operates at Layer 3, which means our common L2 Ethernet switch won't know how to handle the multicast data; it will do one of two things: a) throw it away b) broadcast it to all ports. Both of those are clearly bad . . .

Next month, I will describe how a concept called IGMP snooping can be used to solve the problem . . .

Wayne Howell is the CEO of Artistic Licence, the lighting controls company that he founded in 1988. Wayne invented Art-Net and is actively involved in the ESTA technical standards programme.