

Streaming ACN in anger

I TRY AND KEEP THESE COLUMNS away from war stories as much as I can. On the whole those are only interesting to the participant, not to the audience. This issue however, I'm making an exception as I think the problems I had and the mistakes I made on an installation are useful and might help somebody else from making the same mistakes. The subject is DMX512 with a very large number of universes, and making that work over an Ethernet network in the real world.

“ I was confident that any issues we would come across would be wiring or installation related. We'd probably get some addresses wrong, or have a bad connector somewhere, but nothing more serious than that, nothing systemic. I was wrong. ”

The place is Dallas, a year ago, and I was in the middle of a large installation at the Hunt Building. We were replacing and updating the external display system on the outside of the building which had originally been a very early Versa-Tube installation by Element Labs. The background that's most relevant to the story is that this was a large number of RGB pixels, 53,565 to be exact, which had originally been controlled through a custom video based protocol. Our brief was to replace all the tubes and bring the control system up to date with something that used standard DMX512 protocol. The customer wanted



The façade of the Hunt Building in Dallas.

PHOTOGRAPHS COURTESY OF HUNT OIL COMPANY

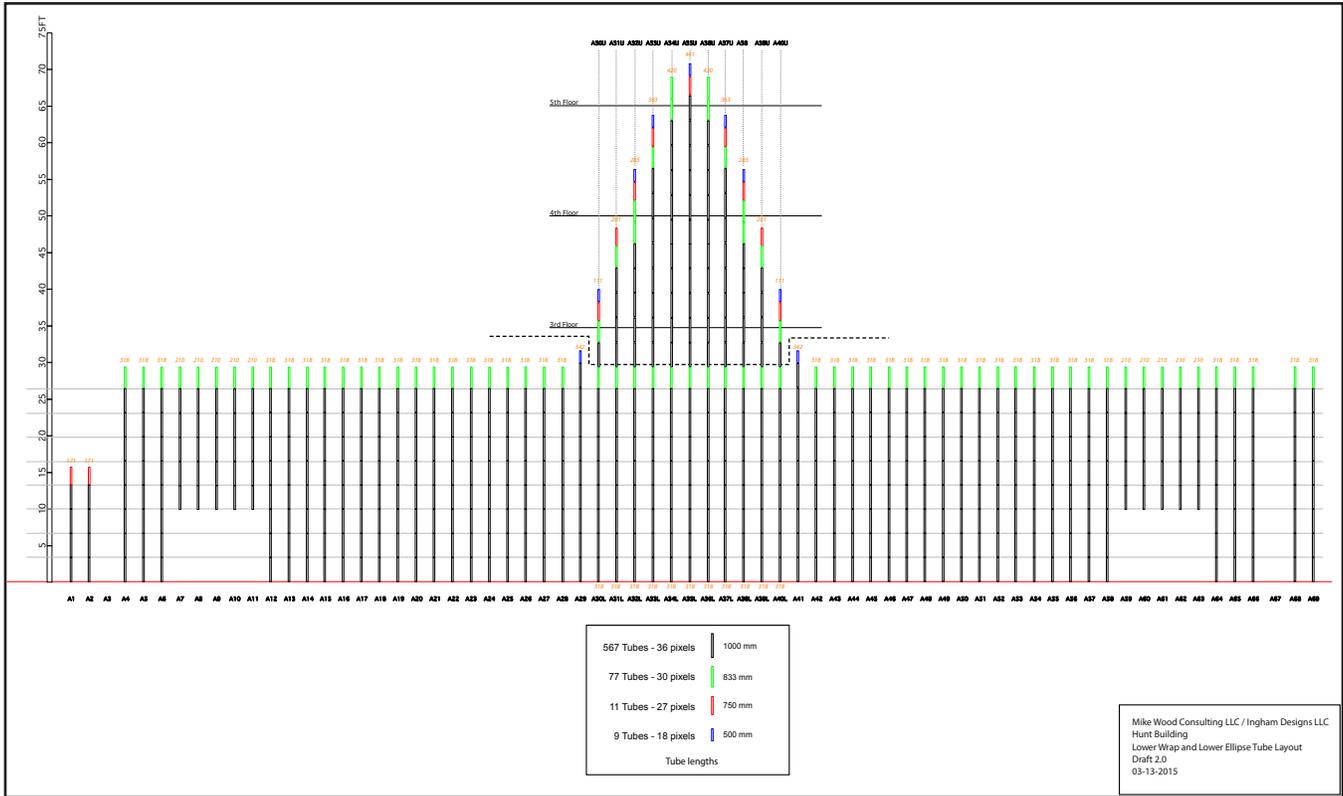
to get everything back to a standard communication method so that they would always have options in the future. This all had to happen without taking the system out of service for more than a weekend at a time.

We did this in two stages, first replacing all the tubes over a period of two years while retaining the original control and protocol. The second stage was to switch the entire system over to an Ethernet-based control over a three-day Labor Day weekend. That's where this story starts.

On the face of it, this shouldn't have been too difficult a problem. Fifty-three

thousand, five-hundred sixty-five RGB 8-bit pixels is 160,695 DMX channels or 314 universes of data. This had to be updated at a 30 Hz video rate.

Although this doesn't seem like a huge number of channels, when we first started this project back in 2013 there wasn't a server on the market that could pixel map that number of channels in real time from video to DMX, although there were a number that came close and had development schedules that promised to get there by the time we needed it. Fortunately for us, the promised R&D schedules at Pharos Controls materialized and, when



Layout drawing for lower 75' of building.

the time came, we had a server, albeit a prototype, to build with. That just left the issue of distributing those 314 universes around the building. Streaming ACN, *ANSI E1.31*, over a dedicated 1 GB Ethernet network was the protocol of choice, and we'd built our receivers with that in mind. (We'd also looked at Art-Net but, at the time, it couldn't handle the number of universes possible on each IP address. I believe that restriction has now gone.) Shouldn't be a problem, lots of capacity on the network. What could possibly go wrong?

In an ideal world, if we were starting this project from scratch we would have taken the single 1 GB feed out of the server straight into a large, high-quality, but also highly-expensive, managed switch made by somebody like Cisco, then distributed multiple smaller subnets around the building to each of the runs of LED tubes. Unfortunately, that wasn't an option. Because this was an upgrade to an existing running building we weren't allowed, nor

would it have been affordable, to completely rewire the building. Instead we had to reuse wiring that was already in place as much as possible with a minimal amount of new cabling.

A few words about the configuration. There was a single server room on the second floor of the building which fed out through three fiber runs to sub-servers on the second, 12th, and 14th floors. That fiber had been carrying DVI video data in the original installation but, luckily for us, we were able to repurpose the same fibers to carry 1 GB Ethernet instead. That got us to three central points from where the data had to fan out to something over 80 distribution points. Each of those distribution points would have one of our 8-universe sACN nodes capable of driving three strands of RGB LEDs. Because we couldn't home-run all these feeds, we were forced into a less than ideal situation where Ethernet was split out in a branching tree through multiple switches. Switch to switch to

switch, expanding each time. The worst runs had four cascaded switches and, in total, we needed around 50 switches to finish the entire installation.

Having to specify 50 switches meant the job couldn't afford \$2,000 switches. Instead they needed to be \$200 at most. Which means consumer or office level switches. That shouldn't be a problem, should it? 1 GB switches are ubiquitous these days, how hard can it be to get 50 off-the-shelf switches and put this together?

“ All I can say for sure is: It was my fault! ”

So, with no hint of the troubles to come, I picked a switch, in the first place it was one manufactured by Mikrotik. You may not know this company, but they make some really excellent, sensibly priced network hardware and I'd had great success with

them in the past. Good prices, and very reliable. I'm not completely naïve; I did buy a single switch and tested it with 30 Hz refresh sACN. It worked fine, so I stopped worrying and got 50 on order.

Fast forward to a month or so later, and Labor Day 2016. We'd spent two days with a crew switching over the entire system from DVI control to Ethernet. This was a one-way conversion. The wiring and firmware changes we made as we went along meant it would not be practical to switch back if we had a problem. I was a little nervous, but not too much; we'd thoroughly tested the system both back in the shop and on a small portion of the building with no problems. I was confident that any issues we would come across would be wiring or installation related. We'd probably get some addresses wrong, or have a bad connector somewhere, but nothing more serious than that, nothing systemic. I was wrong.

Late on the Sunday night we made the last connection and switched the new server on. First tests with static colors looked good. Yes, indeed we did have a few swapped addresses and a bad connection or two, but those were quickly fixed. Then we tried playing full speed video. Hmmm, doesn't look quite right. About one third of the building was absolutely fine, and showed smooth video at full speed. The other two thirds though were slow and jerky, the images were correct, but were only updating at about half a hertz, that is, once every two seconds. What was going on, and why?

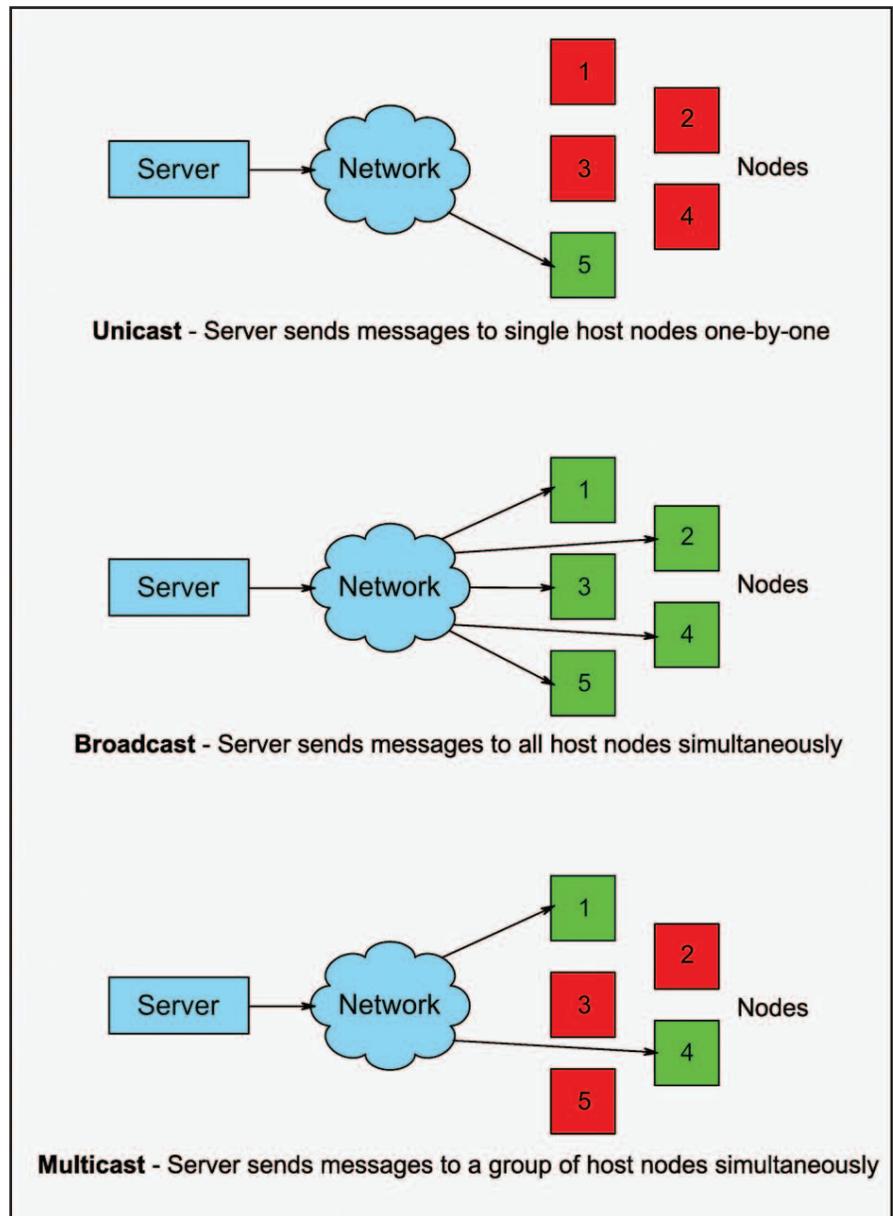
Mild panic ensued. Must be a problem with the server. Must be a bad switch somewhere. Must be a bad cable. Perhaps we are getting noise in the system? Nothing we did made any difference. After sitting down and sketching out what was going on it became clear that the portions of the building that were working were those on the lower universes. In fact, everything up to about universe 60 was working fine, while everything above that was working at slow speed. Why? What could there be about the system that stopped the upper universes from refreshing at the correct rate? Must be

our fault. We had manufactured our own sACN nodes, each of which accepted eight universes of data and converted it to the format to feed each million of LED pixels. We must have done something wrong. We took a node and ran it on its own, testing out both low and high universes. No problems. That wasn't it.

What about multicast snooping? Maybe we were overloading the system and we could improve things by getting snooping running all the way through. Multicast snooping is a method used in routing with supported switches where the switch

examines (snoops) multicast messages and only routes a specific multicast stream through to ports that are actively requesting. It should reduce overall traffic on the network considerably. The switches and router all supported snooping, so it should be possible. An hour or so later and, no, it didn't help, in fact it made things worse in that we lost higher universes completely instead of them just being slow! Turn that off again.

Slowly we started to suspect the switches, maybe we were pushing too much data through them? They were consumer level



after all. Again, we tried some tests, no problem with data rate. Okay, time to get Wireshark out and dig deeper.

To cut what was an all-night adventure short—it turned out that the problem wasn't the data rate, nor was it the ability of the switches to handle large numbers of universes of data. It was all down to the number of multicast groups that each switch could handle.

A note about sACN here. Each universe of DMX512 is carried over a different Ethernet multicast group. This is a clean way to handle the problem as multicast groups are widely supported by standard hardware, and are routable. It also means that multiple receivers can easily look at the same universe if needed without the server having to transmit redundant data. Once the multicast stream for, say, universe 54 is being transmitted, any number of receivers can subscribe to that feed and listen to it without any increase in overhead. Multicast groups are widely used as a way to send digital video data. A channel of television is very similar conceptually to a universe of DMX data. Each has the requirement to be sent to multiple receivers in a one-to-many format and each wants to eliminate redundancy. The difference is that, in any reasonable household, you aren't going to be viewing more than 50 video channels at once, perhaps you have four TVs, each with a DVR capable of recording four channels simultaneously, but that only adds up to 16 channels/universes at a time. We had 314.

It turns out that domestic (and much office grade) networking equipment has a limit to the number of multicast groups it can handle simultaneously. You won't find this information in the datasheet, and, I can tell you from bitter experience, that the support desk for the manufacturer won't know either! Even worse, that limit varies from model to model and version to version of the same switch. Of the 50 apparently identical switches installed on the building,

the 300 we needed. I suspect the number is closely related to the amount of RAM allocated to multicast groups, which explains why apparently identical switches gave differing results. Slightly different firmware builds allocated varying amounts of RAM to multicast. Multicast is likely the bottom of the list for allocation, and just gets whatever RAM is left.

It was now the morning of Labor Day and, although the building was kind-of working, it wasn't working well. We needed to swap out all the switches (or at the very least those above universe 60) for something that would handle the higher numbers of multicast groups. This was a public holiday, so the options were limited. Fortunately, there is a Fry's in Dallas that was open. We headed there and purchased samples of just about every 8-port 1 GB switch they had, including samples from D-Link, Linksys, Netgear, TP-Link, and Trendnet, and took them back to test.

I won't bore you with the testing, but in the end we learned that one of the switches we found that day, a Netgear ProSAFE



Texas Rangers won!

GS108T, would do the job. It turned out it had to be that exact model, and a specific revision number. Other Netgear 8-port switches, although apparently superficially very similar, couldn't handle 300 universes.

Do you have event photos of sACN in use?

To continue the momentum from the cover article of the Summer *Protocol* issue, we are looking for images of shows and/or events that used the ESTA standard *ANSI E1.31*, or sACN.

Protocol would like to create a page or spread illustrating the sACN standard at work. If you would like to contribute, send your photo(s) by December 15 to Beverly.Inglesby@esta.org. Include core caption information: the show or event, approximate

date, the key players, any photo credits required, and note your permission for use. With your approval, we will also include this info on *Protocol's* Facebook page and @ESTA_Standards tweets.

We're starting this project with sACN (*ANSI E1.31*) to highlight how ESTA standards help contribute in many ways to make spectacular, exciting, streamlined, and safer shows and events.

Many could do 200, but not enough. The same for the other manufacturers we tried that day: 100 or 200 was typically the limit, with many handling much fewer. Nowhere in any literature is this limit mentioned. The only way to be sure was to test. Very likely the version of the GS108T you can buy today, a year later, would be different.

Having found a switch that did the trick, we then went back to Fry's, and any other stores we could find open in Dallas, to hunt down as many of the Netgear GS108T switches as we could. We didn't quite find sufficient quantity that day, but, by using the old switches on the lower universes, we had enough to get everything running by the time it went dark and the building looked good. Amazon got us the rest before the next weekend so we could swap everything over to one switch type again.

It's now a year later, and everything on

the building has continued working with—thankfully—very few problems, none of which have been related to either sACN or switches. This was the first time we'd attempted to use sACN, never mind with such a large number of universes. What is pleasing is that all the hardware designed by Ingham Designs worked, sACN worked, and the Pharos server worked. What didn't was the one part I thought was simple: the 1 GB switches. I'm sure if we had gone to more expensive managed switches we wouldn't have had this problem, and the multicast limits would have been documented. But that's the beauty of hindsight. It was a difficult problem to diagnose as we didn't have complete failure, the system worked, but much slower than it should. The switches failed gracefully, presumably relegating the excess multicast groups / sACN universes to a very low priority which

they allowed through when they had the time.

I'd like to say that I'd learned a lesson, but I'm not sure I have. Other than there's no substitute for testing the entire system together. However, with a project of this size done as a retrofit, it's hard to see how we would have achieved that. All I can say for sure is: It was my fault!

Next time this column will return to your regularly scheduled programming, thank you for listening. ■

Mike Wood runs Mike Wood Consulting LLC, which provides consulting support to companies within the entertainment industry on product design, technology strategy, R&D, standards, and Intellectual Property. A 40-year veteran of the entertainment technology industry, Mike is a past President of ESTA and Co-Chair of the Technical Standards Council. Mike can be reached at mike@mikewoodconsulting.com.