Heritage-listed, cavernous and continuously in use, St Pauls presented its challenges.

St Paul's Cathedral

A new distributed, beam-steered speaker system has proven to be an absolute godsend.

— Text: Christopher Holder

Cathedrals: great for a bit of Gregorian chanting, an organ recital or a bell-ringers' convention, but a Synod?

Okay, I'll back up a little bit. For those not baptised into the Anglican church or conversant in five Celtic dialects, a 'synod' is a quaint, olde worlde name for a church committee or board meeting. It's the time for the Anglican high council to get together and talk shop. It also a process much more suited to a modern conference centre rather than a drafty old neo-gothic edifice with a reverb time of about two minutes.

But the decision had been made. And, more than that, Rutledge Engineering had not much more than a month to get in, rip out the old and install the new… in short, they promised a quick-smart, functioning 21st-century audiovisual system before the Archbishop's opening address.

Here was the brief (in brief): a PA for amplifying the spoken word, that would adapt to microphones (wired and wireless) in three main positions around the church. Selecting these positions (as well as some other basic control) would be done by church staff and needed to be foolproof (read 'Overview of the Job' for more).

Anyway… it was a tall order. But Rutledge Project Manager, Harry Michalopoulos, and his team evidently had God on their side. The new distributed PA, wireless control network and analogue I/O were all functional — not fine-tuned or finessed, mind, but functional.

This was thanks in no small part to the fastidious system design of the acoustic consultants, Acoustic Directions. Acoustic Directions went onto fully commission the system, and the results are stunning – the new setup is an unqualified success! Lives have been changed.

AV — more often known for its holier than thou attitude — started off by genuflecting before the high priest of acoustics himself, Acoustic Directions' Glenn Leembruggen.

THE DESIGN

AV: Talk us through what was in St Paul's prior to the refit. Was it up to scratch?

Glenn Leembruggen: It was terrible. It was installed in 1992, yet even then it was already out of date. Why? It was the wrong design. We were asked if the system could be rescued. We said that we could make some improvements to the intelligibility but it would never serve as a solid sound system for the next 15 years. On that advice they wisely decided to start again.

AV: Right. So what were your thought processes when you wandered around St Paul's for the first time.

GL: I was thinking, I would aim to design a sound system that did justice to the architecture, and possessed excellent bandwidth, coverage and clarity.

AV: Very good. But how?

GL: The first decision that needed to be made was whether to use a centralised sound system or a distributed one. Richard Falkinger, the architect engaged by St Paul's, was aware of the, let's call them, hyper-directional speaker arrays, from the likes of EAW, Renkus Heinz, and originally pioneered by Duran Audio [now distributed by Audio Products Group]. These are long horizontal arrays that neatly integrate into the front of the churches architecture. Their acoustic output skims along the surface of the listening plane and they have the advantage of being architecturally simple to integrate — it slips in easily, aesthetically. That was their thinking. But from a high fidelity perspective — and in my book, fidelity is one of the principal ingredients of speech intelligibility — the frequency response is more variable with distance from the loudspeaker [see 'STI — The Leembruggen Addendum' box item]. So I wasn't convinced of the centralised speaker option.

> "Rarely has the transforming power of an audiovisual design been more evident than in St Paul's"

There was another major stumbling block: the church conducts certain parts of the service and various activities in other parts of the church — for example, there's the gospel reading in the service that's conducted halfway up the nave (where most of the congregation is seated). So that means the central array would have been too distant from the person talking — there would be too much of a delay and it would be too difficult to speak.

The other problem was the considerable shadowing from all the pillars. Our coverage zones in the church are not just the nave [or main, central portion of the church] but the other side of the pillars in the aisles as well.

So the combination of those factors sounded the death knell for the centralised speaker system, in this case. It meant that we had to look at an architectural, distributed system to achieve, a) a more consistent frequency response all over the church; b) to ensure the level of the delay, relative to the person talking, would be much lower; and c) to make sure the people behind the pillars in the aisles were being addressed by the sound system.

AV: Can you describe the key principles of a distributed sound system in a big space such as this?

GL: It's all about getting the listeners as close as you can to the loudspeakers and stop sound going into the ceiling space. St Pauls has a high ceiling, so the last thing you want to do is put sound up there. If you do, then you've got the problem of late-arriving reflections. So the idea is to keep the sound directed down at the only acoustically soft 'material' in the cathedral, the people.

It's another thing that is often forgotten about indirect distributed systems: the direct-to-reverberant ratio, that is one of the key factors determining intelligibility, depends on the square of the distance between listener and loudspeaker. The closer you get to the loudspeaker, the greater the direct-to-reverberant ratio, which means significant intelligibility improvements; and that was our design philosophy.

WHICH LOUDSPEAKER?

AV: Right: so you've decided on a distributed loudspeaker system… what next?

GL: The next question was: which loudspeaker? We needed a speaker that embodied three key elements: Firstly, it needed to sit neatly and vertically on the columns, which are at 7m intervals. Secondly, it had to minimise sound being radiated towards the ceiling. Thirdly, it had to steer sound downwards towards the listeners while providing a constant frequency response over its whole coverage area. These are quite difficult to achieve in a single speaker!

AV: Yes, I think I'm detecting a recurring theme here!

GL: The industry is very accustomed to looking at coverage — it's the one plot we most commonly fall back on. But we're not nearly as focused on clicking on the button that tells us what the *frequency response* is at any one listening position. So, my focus is not so much on coverage, it's on consistent frequency response. If we can get consistent coverage — in other words: equal level — well, fantastic, but I'd much rather sacrifice a certain amount of coverage consistency (say, a 3dB variation) for an ultra-consistent frequency response across the entire listening area.

The other big consideration was the budget. The Anglican church does not enjoy the same level of funding of some of the other Christian churches, and therefore our design had to be extremely cost effective.

AV: So, clearly, it's the loudspeakers that have the most influence on the outcome?

GL: That's right, when it comes to these kinds of situations, the loudspeaker is king — it has the capacity to make or break the system. A loudspeaker's radiation pattern and frequency response are the things that really determine what the listener hears. And that's not just who built the speaker and how they built it, it's the adjustments you make; how you set them up.

AV: Which is where Acoustic Technologies comes in?

GL: We've worked with Acoustic Technologies in the past, and we decided to talk to them again. We'd used their ALA07C arrays in the Adelaide

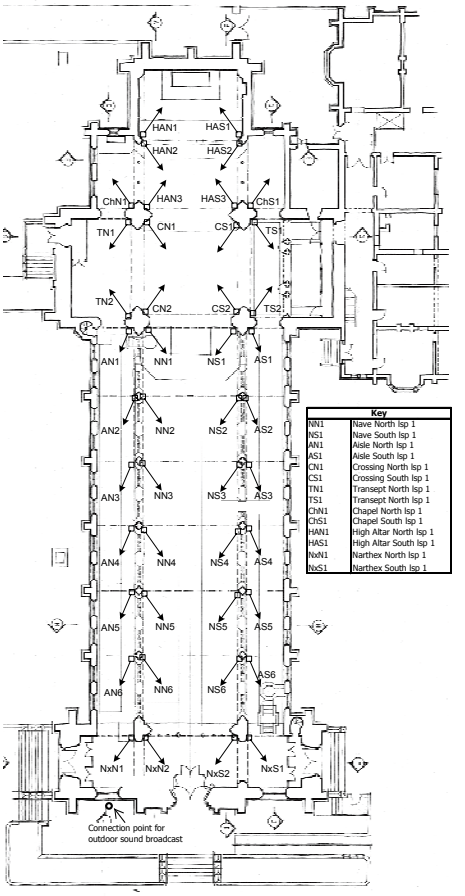

Some 44 Acoustic Technologies ALA07C arrays were carefully positioned on columns throughout the cathedral.

SPEECH INTELLIGIBILITY
THE LEEMBRUGGEN AMENDMENT

Glenn Leembruggen: We're conditioned to think that speech intelligibility is all about scoring a high speech transmission index (STI) figure. Effectively, the STI measures the corruption in speech clarity from background noise, temporal smearing and the ratio of early to late arrivals — in effect, the better your early-to-late ratio the better the speech transmission index.

But the STI is almost completely blind to the effects of the tonal balance of speech. I can say that with authority because I'm part of the international standards committee that is concerned with rewriting the specifications for the management of the STI.

As yet, there is no real way to assess the effect that frequency response has on intelligibility, but I know from experience that it's incredibly important — especially when it comes to listening comfort. Given you have to listen to speakers for prolonged periods in a parliament or a court, listening comfort is really important — it affects your ability to concentrate — so comfort becomes an intelligibility issue. I believe that providing an accurate, or hi-fi, frequency response, where speech is concerned, is just as important as getting good temporal behaviour or a high signal-to-noise ratio. These hyper-directional, horizontal arrays will get through to the listeners, and you can get an adequate temporal behaviour, but you won't necessarily get the desired frequency response.

AV: Why?

GL: The single array is covering the full length of the church, and there's a difficulty in being able to create consistent frequency response from near to far — from 5m to 40m. You can EQ it flat at 20m, but at 40m it would lack bottom end. Up close it would lack top end.

It is one of the fundamental ut even with a two-profile array, it is very difficult to get an even frequency response over a distance range of 3 m to 50 m difficulties with a continuous line source. If you have different profiled boxes, say a J profile for low frequency, and a different kind of J profile for the high frequencies, you can start to get a more consistent response. But even with a two-profile array, it is very difficult to get an even frequency response over a distance range of 3m to 50m

AV: What frequency range are we talking about here?

GL: For high fidelity speech I'm looking at a frequency response from 100Hz to 10kHz. But I'm always trying to reach higher, to 12kHz. There's a little bit of sparkle around 10 to 12k which adds a 'relaxedness' to the sound. People are conditioned into thinking about music being from 20-20k, but not conditioned to think about speech in hi-fi terms, and that extra sparkle does make for a more relaxed listening experience.

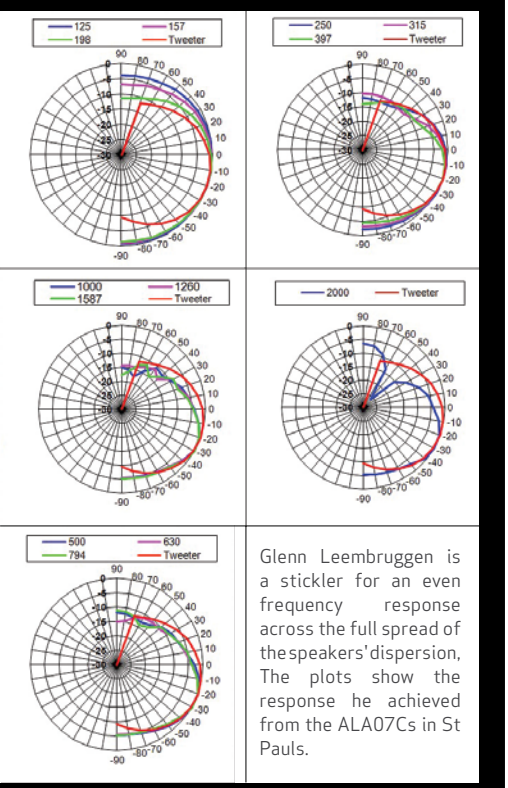

Glenn Leembruggen is a stickler for an even frequency response across the full spread of the speakers' dispersion, The plots show the response he achieved from the ALA07Cs in St Pauls.

law courts and found them to be an effective, inexpensive steered system — regular steered loudspeaker systems are normally thousands and thousands of dollars apiece. The trade-off is that the EAW or Renkus systems are active units and the ALA07C isn't — the amplifiers are in a rack and we have to run cable, as opposed to an active loudspeaker where the amplifier is in the back of the loudspeaker. But, in the 'pro' side of the ledger, the enclosure is slim-line and sits snugly into the fluting of the columns.

AV: What do these arrays look like?

GL: They're eight-channel speakers: seven low/mid frequency drivers, and a high-frequency array of three tweeters. Each array of tweeters has its own passive tapering or shading network to give a constant beam width.

> "I'd much rather sacrifice a certain amount of coverage consistency for an ultra-consistent frequency response across the entire listening area"

The goal of the array is to maintain a constant radiation pattern across all frequencies, which will mean that wherever the listener is, they'll get a constant frequency response. But as well as that, we tailored the radiation pattern to compensate for the distance loss — the change in frequency response from 1m at its closest, out to the maximum beam distance of 7m [after 7m the next array would take over].

AV: Which is achieved via the DSP processing, presumably?

GL: Yes, the radiation pattern and frequency responses of the low/mid drivers are controlled by our active signal processing algorithms, while those of the high frequencies are controlled by the passive network that is associated with the tweeter array, which is physically tilted. Remember: it's almost impossible to electronically steer high frequencies, you simply can't space the tweeters close enough, so it was a case of manually pointing the tweeters in the right direction.

AV: Can you tell us more about the processors?

GL: The processing was also from Acoustic Technologies. Through the course of our work on the Adelaide law courts and St Paul's, they developed an active processor and steering box of electronics. Previously, we used off-the-shelf DSP from Biamp — the Nexia — to do all the signal processing and steer the loudspeaker. But in this particular instance we decided to use the Acoustic Technologies product. It's called the FLA8 and apart from multi-channel amplification, it holds the beam-forming parameters and crossovers that we developed.

INTIMATE SOUND

AV: You could probably describe the St Paul's system as 'intimate' — like someone whispering in your ear or tapping you on the shoulder. Agreed?

GL: Right, I agree. And that's a function of a good early-to-late ratio, and a demonstration of the fact the design was set up with care and accuracy. We put in two weeks of commissioning, checking and confirming every setting — with 44 arrays and 25 amplifier/processors that's a serious number of connections that need checking, but we made sure every one was correct right from input to acoustic output. And then there's the setting up the equalization, level adjustments, delay…

AV: So, the system is time-aligned as you sit further from the pulpit?

GL: Yes, the loudspeaker at the front [behind the pulpit] is Time Zero, and then progressively delayed. Or, in the case of the gospel reading, half way down the church, that position becomes Time Zero. The delayed nature of the system contributes to that sense of intimacy — it pulls you right in. At one point I recall being in church when the Archbishop was speaking right at the front of the church at the high altar. He was using one of the Countryman mics — a very nice headset — and I was stunned. He was right 'there'. It was then that I thought: we've done a very nice job.

AV: Can you tell me about the system presets?

GL: Each setting on the AMX touchpad activates the microphones and speakers relevant to that preset. There are a couple of presets where all speakers aren't on, but the average service has almost all microphones and on almost all speakers on. We rely on the auto mixer to only turn on those mics that are being spoken into. It's almost a set and forget system. [More on the presets and the control system later.]

AV: What's taking care of the auto mixing?

GL: Again, that's the Biamp Audias. We think they're a good deal more intelligent than the other auto mixers on the market. That's the reason why Biamp has been our primary choice of DSP for speech reinforcement applications for such a long time now. What's so special about its auto mixing? Apart from sensing arrival times, it has an interesting algorithm that looks at the wave shape to determine whether a mic is active or not. It's one of our key ingredients to achieving high gain before feedback and contributes to that intimacy we were talking about.

NOT JUST ANOTHER CATHEDRAL

Rutledge Engineering is no stranger to this type of job — having taken care of St Mary's in Sydney, it knows its apse from its elbow, and it slipped into this delicate role with all the discretion of a troupe of 'on best behaviour' altar boys. Harry Michalopoulos — Greek Orthodox, would be my guess — was the Project Manager.

AV: How does St Paul's compare to previous cathedral jobs?

Harry Michalopolous: This job has more locations than any other we've done previously — more speakers, patch points etc. But giving a cathedral crystal clear and intelligible sound requires a special solution.

AV: Sounds like there were some complicating factors.

HM: Yes, but fortunately we didn't have to rough-in the cable. That was Rick Clark and his team from Rely On, who have been the cathedral electricians for many years. When we arrived, it was a case of: thank goodness someone else is roughing this cable in! They have an intimate knowledge of the cathedral, and that coupled with the torturous heights involved — the access issues — we counted ourselves lucky.

The heritage nature of the cathedral also meant we had to be especially careful — there was a heightened sense of sensitivity to the surroundings. For example, the pews could be moved, but around certain columns there were other furnishings that we just had to deal with. So certain locations required — I wouldn't say 'extraordinary' efforts — but some special solutions.

AV: How did you fit the speakers to the columns exactly?

HM: Glenn had done precise measurements. Not just on the array heights and angles, but of the geometry of each column — the fluting of the column. The speaker had to be mounted precisely in a particular flute. We had to put a threaded rod in to these intersections, but make sure that it would stick out at the correct angle so we could actually bolt the array's brackets there.

We took the problem to our Chief Audio

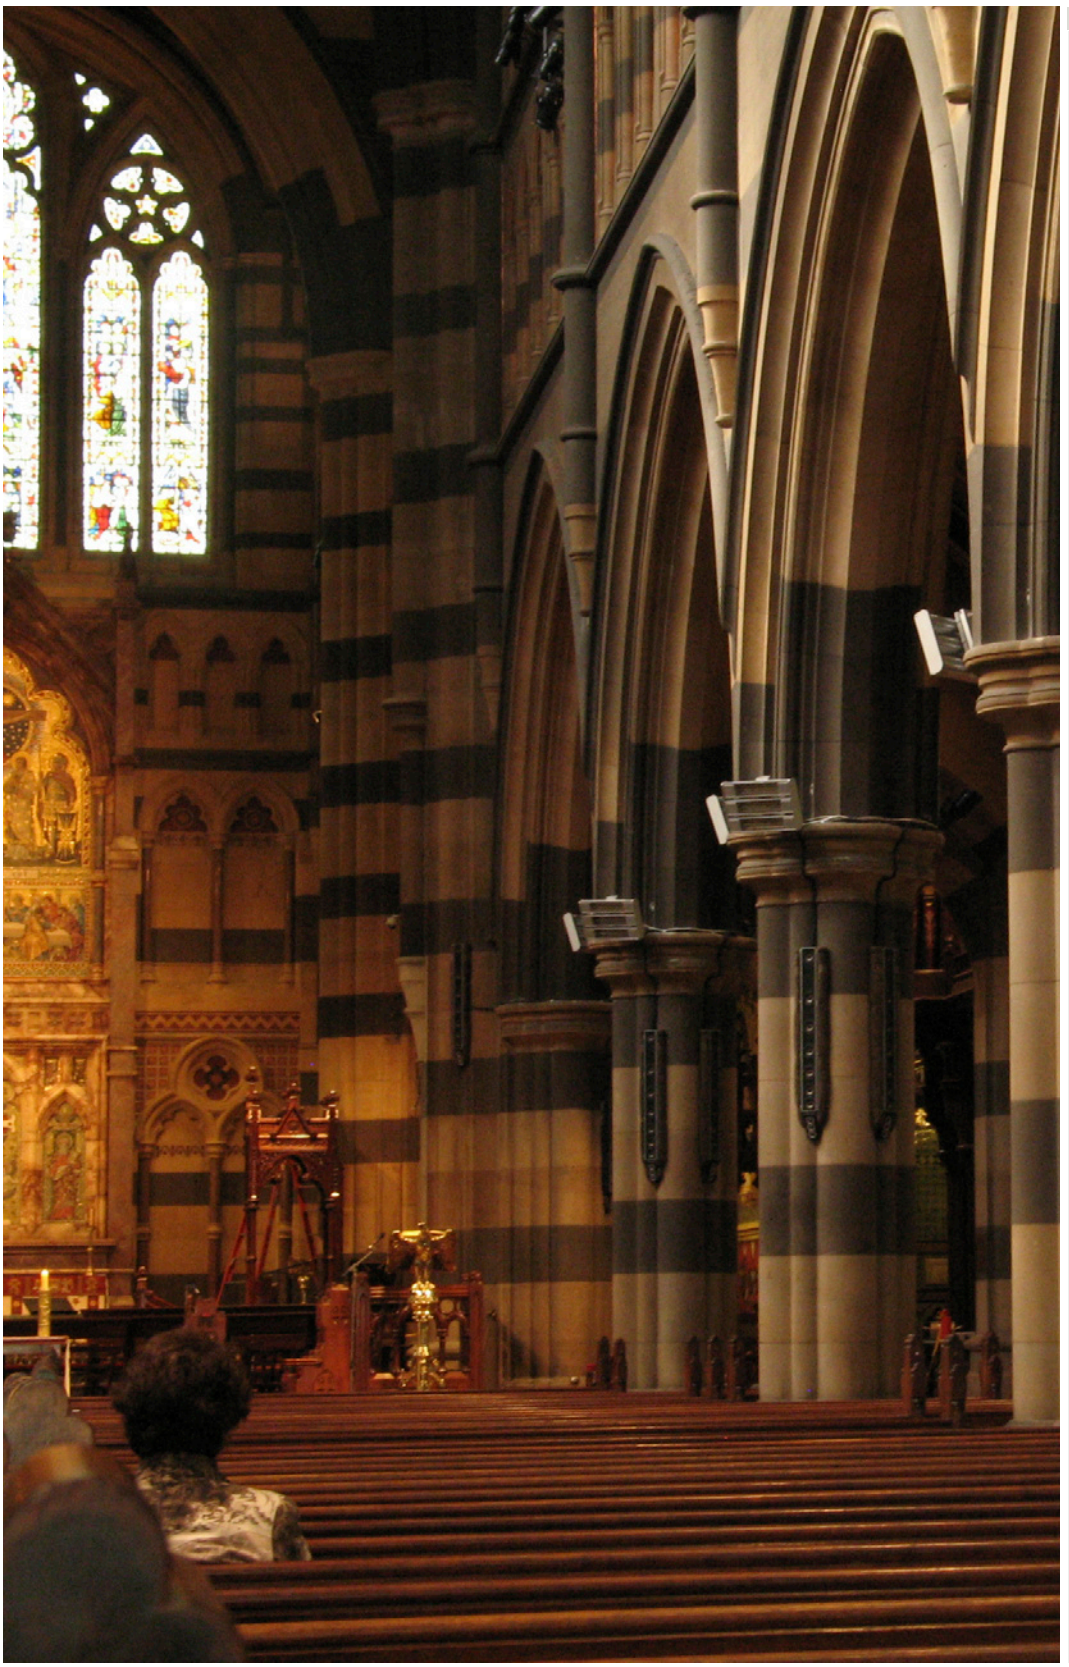

OVERVIEW OF THE JOB

• The main aspect of the brief was to provide a crystal-clear speech reinforcement system. Then there was the multitude of analogue I/O to all the microphones and touchpanel points around the cathedral, and the wireless LAN.

• The system was based on 44 modified ALA07C column speakers from Acoustic Technologies. Each box comprises seven low/mid drivers and a tweeter array on the baffle.

• The bulk of the church — the nave and aisles — were covered by two ALA07Cs per column, all mounted at the same height. Electronically beam steering the column speakers allows the sound to be angled down to a reference axis — the ears of the majority of the seated congregation.

• Adjacent speakers on a column are acoustically decoupled using all-pass filters. to mitigate the effects of phase interference.

• The ALA07Cs are a passive speaker. The amplifiers are in racks. Acoustic Technologies has also developed the DSP processing required to take care of the beam steering. The DSP has been integrated into the eight-channel amplifier chassis. The resulting device is called the FLA8.

• 25 FLA8s are employed. The church is symmetrical, and thus pairs of ALA07Cs often share the one FLA8 amp/processor. The remaining (non symmetrical) speakers get their own. Each amplifier/processor was precisely married to a particular speaker(s).

• The tweeter array is driven by the eighth channel of the FLA8, and has a passive network to shape its response. But you can't electronically beam steer the HF section, so it needs to be mechanically pointed down.

• The other DSP processing — signal limiting and cancellation, EQ, system delay and auto mixing — is managed by Biamp Audiaflex processors.

• Crucial to the success of the system was ease of control. The control is based on two AMX 8400 roving touchpanels (combined with the Cisco Aironet WAP).

• Rounding off the audiovisual refit is the analogue I/O connectivity: the patch points and multi-core links to ensure the cathedral works well as a performance space.

"It looks like a fairly sophisticated, bespoke installation but is, in fact, a cost-effective, almost entirely off-the-shelf solution. If we had to use commercial, self-steered product we would have paid double the price." – Glenn Leembruggen.

A shot of the Acoustic Technologies arrays in position. The sound is electronically steered down to the ear-level of the congregation.

ACOUSTIC DIRECTIONS TEAM

Acoustic design & commissioning: Glenn Leembruggen & David Gilfillan
Project Management for St Paul's & Acoustic Directions: Mark Hanson

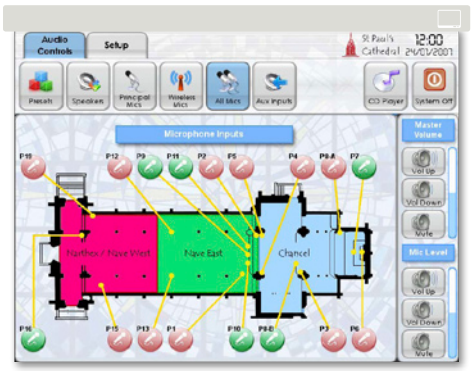

The AMX system control provides easy access to a selection of presets, as well as the ability to turn on/off microphones, speakers, playback and additional auxiliary inputs.

Engineer, Bill Tauscke, who's been working at Rutledge for I don't know how many years. Bill designed this fantastic jig (that will never ever be used again, because it is only relevant to this installation). We had two of these jigs, mirror images of each other, which matched the profile of the pillar. We ratchet-strapped on the jig, drilled straight in, put the threaded rod in, and mounted the speakers. Bill's jig was a big success.

CH: Can you elaborate on this 'crystal clear sound' you mentioned earlier?

HM: It's a speech reinforcement system. It's meant to provide an unplugged sound — acoustic invisibility, if you like. When an untrained person says that the speaker nearest him wasn't on — despite the fact he was hearing everything clearly — that is the ultimate compliment.

CH: A big part of Rutledge's brief was to make the *system* all-but invisible to the church, which included designing a fool-proof interface.

HM: Absolutely. Obviously all the crucial system settings were safe, with the users locked out, but aspects like certain mic gains, system level, turning certain speakers off/on depending on how many people are in the church — all needed to be easy to control.

CH: Based on AMX touch panels, I believe?

HM: That's right. The cathedral decided on two wireless touch panels — the AMX 8400 8.4-inch model. Then, rather than opting for the standard wireless access points (WAPs), we opted for the Cisco Aironet range of WAPs. We've got a WAP at either end of the cathedral, and we've added a third WAP in the administrative corridor, which is otherwise a real dead zone — don't expect to get mobile reception there.

AV: Can you talk me through how the interface works?

HM: It's based on three main presets, which are activated depending on what type of sermon, message or reading is being delivered. Each preset applies some equalisation that takes into account the architecture surrounding the microphone. For example, a headset microphone might be roving, in a particular area. So there is equalisation, level controls, muting or ducking of certain speakers… multiple factors that come together to deliver the best sound for that location.

SAINTS PRESERVE US

My first meeting for this story was on site in the cathedral's Canons' vestry. Seeing Harry Michalopoulos chatting amiably with St Paul's Colin Reilly and Rachel McDougall, it occurred to me what an accomplishment this was. Every job requires an empathy — an understanding of the business or activity being pursued — and a cathedral's requires a little more effort to understand than most.

Because, let's not forget that there's a spiritual dimension at work here. Or, if you have trouble with the word spiritual, then 'emotional'. Suddenly Rachel MacDougall is able to more easily 'lose' herself in the singing of evensong, and the 'PA' becomes transparent to the congregation. Overnight, the Archbishop is able to 'speak' individually to parishioners from some 50m away at the high altar. Rarely has the transforming power of an audiovisual design been more evident than in St Paul's.