Birds on the wire – Olivier Messiaen's Livre du Saint Sacrément in the world's first wave field synthesis live transmission

Jörn Nettingsmeier
Lortzingstr. 11
45128 Essen
<nettings@stackingdwarves.net>

On July 29 2008, a world premiere took place between Cologne Cathedral and Technische Universität Berlin (TU): the first live concert transmission using the wavefield synthesis technique.

Organist Prof. Winfried Bönig performed Olivier Messiaen's "Livre du Saint Sacrément" on the two Klais Organs of the cathedral for an audience of around 1500. At the same time, 600 people in Berlin listened to a live reconstruction of the concert on the wavefield synthesis array at the TU, supervised by Tonmeister Prof. Dr. Stefan Weinzierl.

This is the project report for the Cologne side of the production.

1 Planning

The idea of a WFS live transmission was originally devised by Prof. Stefan Weinzierl of TU Berlin and Dom organist Prof. Winfried Bönig.

The goal was to capture a convincing representation of the room ambiance of Cologne Cathedral as well as near-field images of the organ works and remote stops, to allow for a spatially correct re-creation of the live concert experience in the WFS auditorium of the Technische Universität Berlin (TU). The signals were to be transmitted to Berlin in realtime, and pristine audio quality.

In May 2008, the author was approached by the TU to provide the necessary network infrastructure and to act as project manager for the Cologne side of the production.

The sound engineering work was to be taken care of by Tonmeister Hans Schlosser of the Robert-Schumann-Hochschule in Düsseldorf (RSH), and a team of his students. A sound truck and all necessary audio equipment would be kindly provided by RSH.

With the help of WFS specialist Dipl.-Ing. Frank Melchior (of the Fraunhofer Institute for Digital Media Technology IDMT, Ilmenau), we had established that up to 24 discrete signals of uncompressed audio would be desirable for the transmission.
It was decided early on that the actual transmission should happen via an Audio-over-IP link, possibly based on UDP/RTP.
At a sampling rate of 48 kHz and a word length of 32 bits, the data adds up to a bandwidth of roughly 37 Mbit/s. There exists a fiber optics link from the *Silberkammer* in the

Illustration 1: The "Querhausorgel", the main organ of Cologne Cathedral (Klais, Bonn, 1948, 1956, 2002)

Dom basement to the Domforum and the Dombauverwaltung (adjacent buildings across the Roncalli square), but the available uplink is a meagre 4 Mbit/s ADSL line.

Rental of one or more additional DSL lines was not feasible, since the local ISP had no provisions for line trunking, which would have been undesirable anyway due to the increased possibility of out-of-order packets. Moreover, their quotation for only one temporary DSL link at 18 Mbit/s came close to exceeding our total budget.

So we decided to cooperate with the

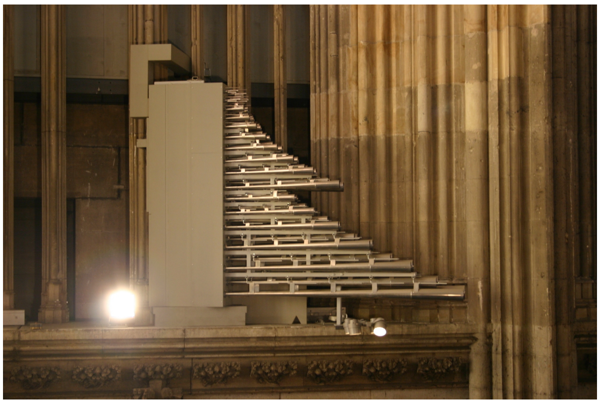

Illustration 2: One of the two high-pressure tuba stops on the Triforium, above the main entrance (Klais, 2006)

Illustration 3: The "Schwalbennestorgel" on the northern wall of the central nave (Klais, Bonn, 1998). Mechanical and (optional) electric action for remote control from the main organ console.

Kunsthochschule für Medien (KHM) (which had hosted the Linux Audio Conference in 2008, and whose staff the author had worked with before) to try a point-to-point optical link to the Dom via infrared laser. The KHM has direct line-of-sight to the Dom spires at a distance of about 850 m. For the uplink, the Citizen Media research project at KHM kindly let us share their 2 Gbit/s connection.

After many failed attempts to obtain a suitable transmission system from several large national rental companies, OPTEL, a laser link vendor from Hamburg, came to the rescue. Not usually in the rental business, they found the project sufficiently challenging and interesting to provide us with a pair of spare laser devices and extensive technical support for the duration of the project.

About two months before the concert, the author contacted the Dombauhütte to ask for a permission to install the laser and run a fiber cable down the spire to the planned location of the sound truck, and to determine ways to mount microphone arrays in the Dom itself. Dombaumeisterin Prof. Dr. Schock-Werner called a meeting with some of her senior staff members a few days later, where we presented the project. We won their support, quickly arrived at a consensus and developed a first concept for microphone placement and cable runs with the assistance of the technical staff of the Dom.
Now the real work could begin.

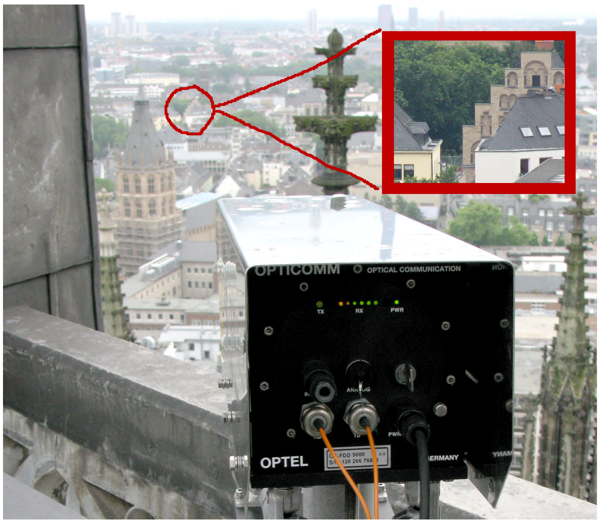

Illustration 4: The laser link seen from the intersection spire. The other end was located to the left of the staggered facade (the Overstolzenhaus, KHM library building).

1.1 Microphone distribution

Tonmeister Hans Schlosser and his students checked out the location one week before the concert, to estimate required cable lengths and to decide on the distribution of the remote mic preamps and the final microphone placement.
In total, we used 20 discrete signals. Five Schoeps MK 5 omnis in two AB pairs and an additional mono spot mike had to be used on the main organ due to its complicated layout, while an extra M/S pair (Schoeps MK 5/MK 8) covered its *Rückpositiv.*
One sub-cardioid AB pair (Schoeps MK 21) was facing the new *Schwalbennestorgel* ("swallow's nest", since it clings to the wall suspended from the roof on four steel cables) from across the main nave.
Two mono hypercardioids (Schoeps CCM41) were aimed at the high-pressure *tuba* stops over the entrance.
All these signals were to be rendered on the WFS array as point sources at the appropriate locations.

In addition, eight figures-of-eight (Sennheiser MKH 800) were used in a Hamasaki-style array to capture uncorrelated ambiance information (four were oriented towards the side walls and four to the ceiling); they would later be rendered as plane waves arriving from the four corners of the WFS array.
Finally, one Sennheiser MD421 covered the announcer.

During their visit, the RSH team decided that the ambient Hamasaki square should be relocated from its originally planned position in the southern transept to the apsis of the central nave, to reduce the amount of direct sound from both organs.

2 Network installation

Two weeks before the show, the author arrived at the Dom early in the morning to run an optical cable down from the *Vierungsturm* (the intersection spire, at the crossing of transept and central nave) at around 70 metres to the ground, which was to connect the stream server to the laser link.

Illustration 5: Voltage measurement at the initial testing location of the stream server: the "bunker", an appendix to the southern transept used as a storage closet.

Rope access techniques had to be used in the upper section to provide sufficient anchoring for high wind conditions, but rest of the cable run was easily accessible by staircases and via fixed scaffolds. With the help of Head of Scaffolding Wolfgang Schmitz , the cable was secured in place in less than three hours.

Chief electrician Lothar Reinhardt had already taken care of running a power supply to the roof, and around noon everything was ready for the installation of the lasers. The initial setup was performed by two skilled OPTEL technicians, and the laser link was up in less than 4 hours.
Meanwhile, KHM administrator Bob O'Kane had patched our fibre through to the internet.
It may sound a bit childish, but the very first ping from the Dom's gloomy storage “bunker” to Berlin was quite an experience!

2.1 Power issues

However, after half a day of stable performance, the link died.
Checks on the next day revealed that a media converter at the Dom was getting stuck regularly.

Some quickly deployed watchdog scripts showed a regular outage pattern: the link would die in the mornings, some time between 8 and 11 hours, and come back at night, between 23 and 3 hours.

It turned out that the voltage of the the power grid at the Dom was 218 rather than the standard 230 V, and it would dip some more when the load was high. This would cause our media converter (which used a cheap wall-wart type power supply) to choke frequently. Finally, after one week of heroic compensation efforts, an on-board voltage regulator shed its magic smoke and died.

OK, time for another media converter. Luckily, the new one ran on 12 V and so we could feed it from the 12 volts rail of the encoding PC, whose switching power supply had no trouble compensating for the low grid voltage.

2.2 Botanics

Now, just one intermittent fault in a network that we only had remote access to most of the time[1] would have been way too easy.
To keep life interesting, we began to observe mysterious packet losses a few days after the the power issues had been taken care of, occurring in bursts of several seconds duration. After ruling out network and power issues on both ends (which took quite some time) it was found that a tree in between the KHM and the Dom had taken advantage of ten days of very wet weather and three days of sunshine and had been sprouting at an amazing rate. We had not anticipated this to be a problem, because the tree stood slightly

1 On-site access at the Dom was strictly limited to work hours, in this case from 6:30h to 15:30h, and would require a one-hour train ride.

to the side of the laser's line-of-sight, but it turned out that even a light wind would bend the soft new branches into the way, sometimes occluding the beam for more than three seconds at a time. Obviously, this would be fatal for our intended usage.
Thus, it was back to square one: the laser at the KHM had to be raised by at least one metre for the beam to stay clear of the branches. Since the KHM library building we had installed it on is a protected historic site (known as the *"Overstolzenhaus"*, former town estate of a Cologne merchant family), there was no chance of drilling holes into walls, so we had to rely on a couple of stone plates to weigh down our makeshift laser scaffold, made from hastily purchased 60x5mm aluminum tube.
OPTEL shipped us a pair of cross-hair scopes and some measuring equipment for proper alignment via express at a day's notice, and we got to play with lasers again. With some telephone support by Optel CEO Ralf Magiera, the link was re-established a few hours later.

Illustration 6: Keeping track of the microphone setup. To the left, the Yamaha DM-1000 console used for monitoring and signal routing.

Illustration 7: One of the RME Micstasy remote controlled preamps, secured to a column on the Triforium. In the background, the Hamasaki frame can be seen.

2.3 Transmission tuning

On the weekend before the show, the author went to Berlin to run some final performance tests with WFS cluster admin Thilo Koch. After solving some initial buffering issues and local digital clocking problems, we finally found a dependable setup and enjoyed 24 channels of uncompressed 16 bit test signal streamed live from Cologne Cathedral. At long last, we were beginning to relax.

Unfortunately, there was a minor flaw in our reasoning: a network that performs well on a Sunday need not necessarily behave on a weekday.

3 Rehearsal day

On the day before the show, the Cologne crew arrived at the cathedral around 10:00h. The sound truck got caught in a traffic jam, but after the midday prayer at 12:30, the setup work could begin.

The RSH students got busy running a MADI fibre from the sound truck to four remote-controllable RME Micstasy microphone preamps which were placed on the *Triforium* (the Dom's inner balcony), next to the main organ and across from the *„Schwalbennest“* organ. The stream server (which had sat in a closet during the previous testing phase) was relocated to the sound truck and started transmitting while the microphones were brought online one after the other.

Illustration 9: The RSH sound truck.

After some listening, the students decided to swap some capsules for narrower directivity patterns to reduce the amount of reverberation picked up by the spot mikes.

The Dom scaffolders had rigged a radio-controlled chain hoist for us which was used to suspend the Hamasaki square. Two lengths

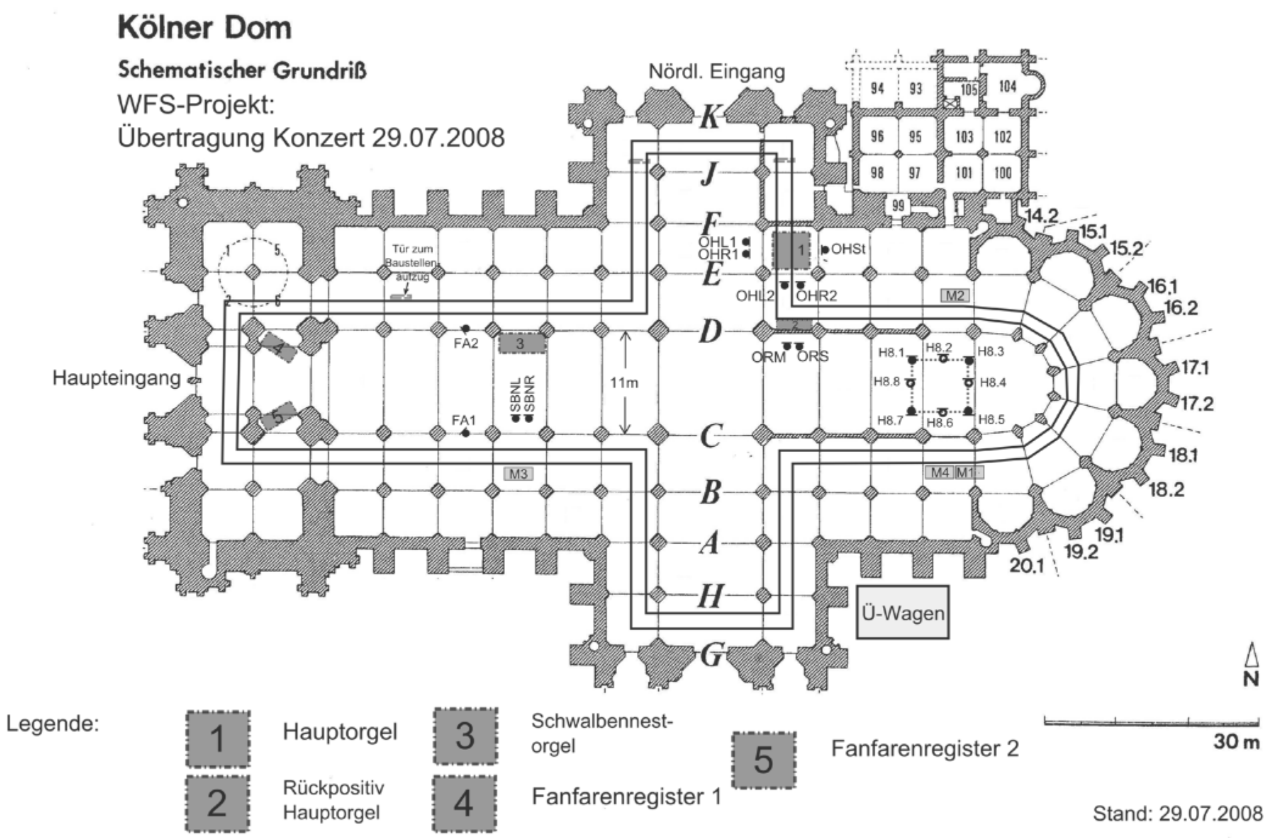

Illustration 8: Overview of microphone positions

of fishing line were then strung to the *Triforium* for precise alignment.

During the day, Frank Melchior and Hans Schlosser took care of an overwhelming number of interview requests by press, radio and TV journalists. Their initiative gave the rest of the team time to concentrate on their work, and as a welcome side effect considerably enriched the media coverage with sound technical expertise.

A video camera with software encoder was set up on the Triforium to provide live imaging to the remote audience in Berlin.

At 17:00 hours, the Berlin crew finally had access to the WFS auditorium (it had been blocked by regular lectures during the day), and began setting up the receiving machine. After some digital clocking issues had been resolved, the first test signals were routed to the WFS array.

3.1 Network problems

We quickly found that the transmission quality was abysmal, with at least one drop-out per second in one of the channels, and frequent losses of connection. The level of relaxation dropped accordingly.
Measurements from several remote networks established that contrary to all expectations, the laser was by far the fastest and most reliable link in the transmission chain, and that UDP packet losses on the order of 10-20% appeared to occur on the campus network in Berlin.
It was decided to get a team together to create a dedicated link to the receiving machine that would circumvent possible bottlenecks, but since it required the cooperation of network operations people at the TU, it became clear that this could only happen the next morning.
At the same time, we found out the hard way

Illustration 10: Detail of the Hamasaki square, with one of the horizontal figures-of-eight (a Sennheiser MKH 800). The array was suspended in the central nave close to the apsis, at a height of around 20 metres.

that all the little knobs that our software provided to deal with network problems (input buffering, redundant encoding, output buffering), either did not have any significant effect or were not working at all at the high channel counts we were using.
Thus, the organ rehearsal which began at 20:00h could not be used to optimize the WFS rendering and was instead spent debugging network issues, which unfortunately remained unresolved.
At around 23:00h, the RSH crew left the Cologne site after having taken a number of impulse responses from the cathedral, while the author stayed behind in the sound truck trying to come up with a plan B.

4 Show time

Work resumed at 09:00h the next morning. While the Berlin team were hoping for a response from campus network operations, the author set up a 16 channel Ogg Vorbis encoder with 160 kbit/s per signal as an alternative transmission system, in the hope that its HTTP-based transport layer would be able to compensate for packet losses at the cost of increased latency (which did not matter for our intended application). The Vorbis stream was relayed via an off-site Icecast server that had originally been intended to serve a stereo-only stream over a separate DSL line as a fallback solution in case the main uplink went down.

At this point, we had one 24-channel PCM stream and one 16 channel Vorbis stream going out, but both could not be received properly at the TU end: the PCM stream due to continuing network problems, and the Vorbis one due to problems with multichannel support in the decoders we used. A special team was dispatched in Berlin to work on a solution. Finally, CCRMA wizard Fernando Pablo Lopez-Lezcano, then guest professor at the TU, came to the rescue and found a reliable way to decode the 16 channel Vorbis stream, using his private laptop to save the show.

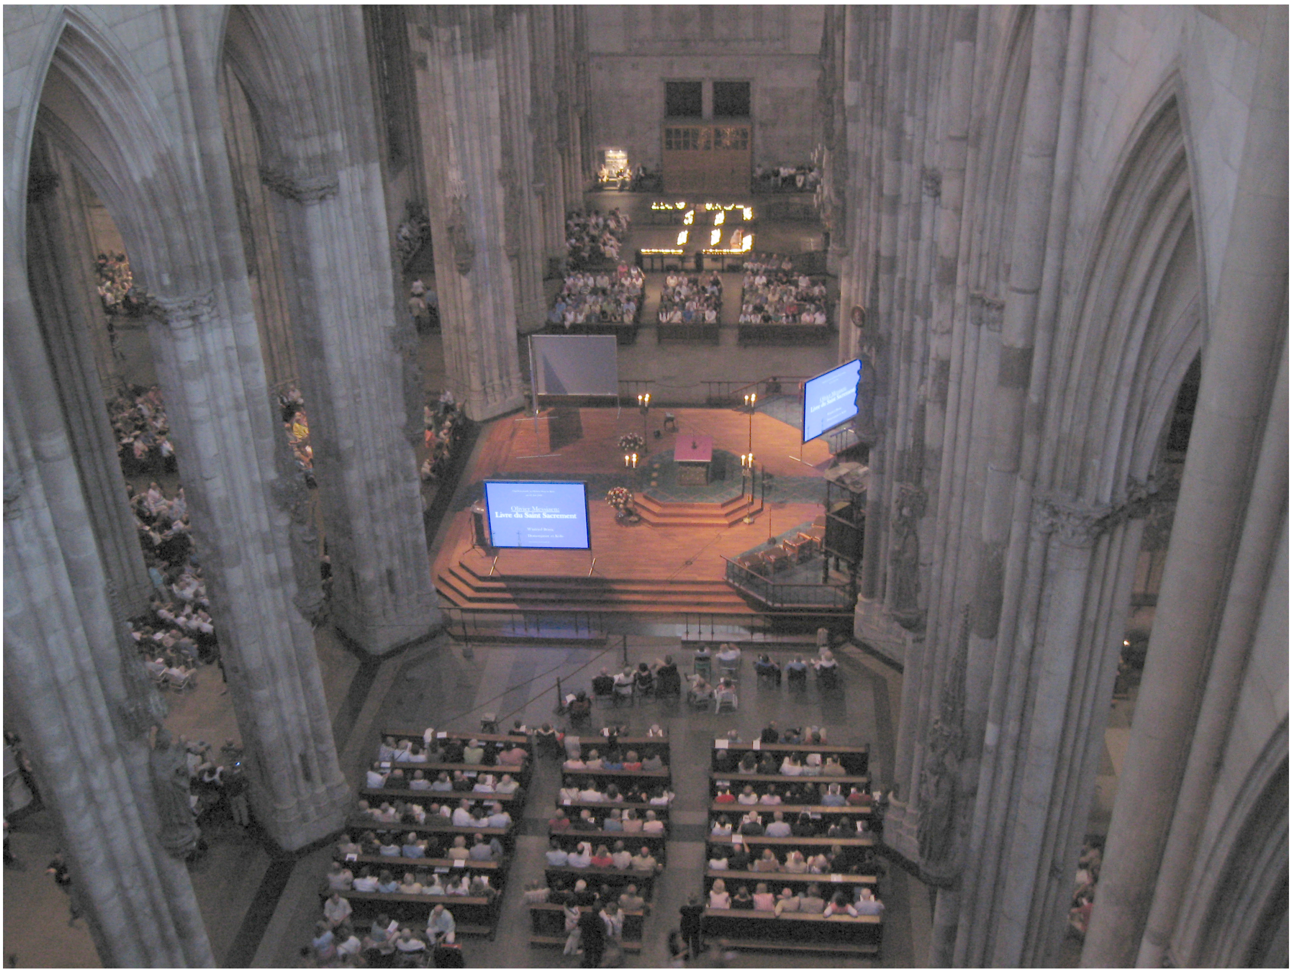

Illustration 11: During the concert. View from the triforium in the southern part of the transept. The video screens around the altar display textual comments to each movement, as intended by the composer.

In the early afternoon, we had a dependable link at long last. The evening mass was transmitted as a test case and used by the WFS operators to align the virtual sources in Berlin. Unfortunately, not all organ works were in use during the mass, so the last few sources had to be fine-tuned after the beginning of the concert.

The concert itself was transmitted without major problems, although we had three dropouts, one during the introductory address, a very unfortunate one in the middle of the performance and one during the final applause. We later traced them to a configuration error of the Icecast stream server. It had been optimized for use as a two-channel fallback, and its client buffer (i.e. its lenience towards clock drift or network issues on the client side) had not been increased to allow for the higher channel count. Effectively, this resulted in only a few seconds of drift tolerance, after which the client would be kicked off.

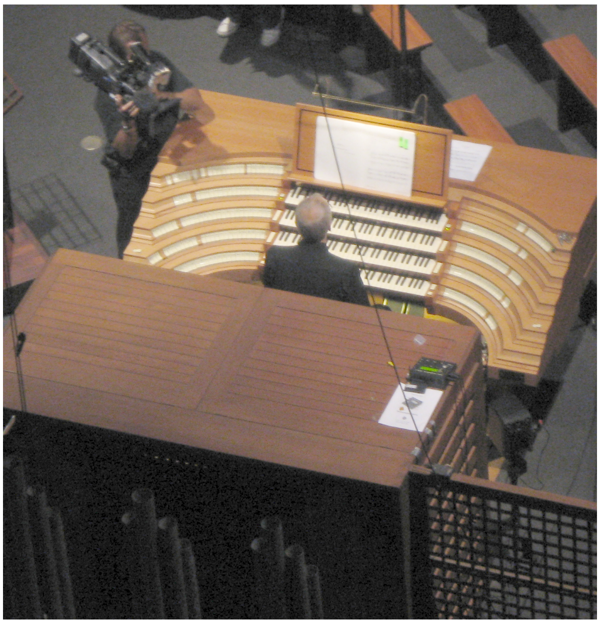

Illustration 12: Dom organist Prof. Winfried Bönig at the console of the main organ, which also provides access to all stops of the Schwalbennestorgel. In the foreground, part the organ's Rückpositiv is visible.

Illustration 13: Tear-down after the show.

Unfortunately, the nature of the problem was not immediately recognized, and we wrongly decided to bring down the video stream to rule out network congestion as a reason for the dropouts. Fortunately, the video service could be re-established quickly after the Berlin crew had rigged a dedicated wireless link.

After the show, the RSH crew began to tear down the equipment. As soon as the public had left the venue, the Hamasaki square was brought down and dismantled. Meanwhile, the recordings (which amounted to around 24GB) were uploaded to a backup location at the KHM via the laser link, for later research an possibly another public replay on the WFS array in Berlin. For safety reasons, the author decided to postpone the the rope work on the roof to the following day, to be tackled after a good night's sleep.

The day after the show, the fibre to the spire was reeled in, the lasers were taken down and shipped back to OPTEL and various pieces of borrowed computing equipment were returned. At 18:00 hours, the work in Cologne was completed.

5 Conclusion

Since the Berlin audience had been primed as to the highly experimental nature of the whole undertaking and the technical challenges involved, the brief outages during the performance were received sympathetically. The sonic result was generally well appreciated and enjoyed by the concert attendants in the WFS auditorium, although the material itself proved challenging for parts of the audience on both sides.

The different parties involved in the Cologne side of the production cooperated extremely well, and similar future undertakings can be expected to be met with continued interest and support.

6 Appendix

6.1 Credits

Special thanks (in alphabetical order) to:

- **Juan-Pablo Caceres,** composer, JackTrip co-developer and lecturer at CCRMA, for lots of online support (even while he was on vacation)
- **Karl Heyes**, Icecast/IceS developer, for remote support in setting up a 16-channel Vorbis encoder
- **Bob O'Kane**, KHM network administrator, for being really helpful at really short notice, network support and lots of little hardware gadgets to save the day
- **Ralf Magiera** at OPTEL, Hamburg, for providing the laser hardware, and for their enthusiastic support
- **Wolfgang Schmitz**, head scaffolder at the Dombauhütte, and his team, for helping to run the fibre link down the spire, and for providing a chain hoist for the ambient microphone array
- **Prof. Dr. Barbara Schock-Werner**, Master Builder of Cologne Cathedral, for her generous and professional support
- **Dr.-Ing. Thomas Schumacher**, architect at the Dombauverwaltung, for interesting insights and lots of overtime on our behalf
- **Lothar Reinhardt**, head electrician (Dombauhütte), and his team, for power supply and countless little favors
- **Martin Rumori**, research associate at KHM, for organisational support and a helping hand in tight situations

6.1.1 Project Partners:

- **Dommusik**, the musical department of Cologne Cathedral
- **Institut für Audiokommunikation** at Technische Universität Berlin
- **Institut für Musik und Medien** at the Robert-Schumann Hochschule (RSH), Düsseldorf
- **Kunsthochschule für Medien** (KHM), Köln
- **OPTEL** Optical Communication, Hamburg
- **CITIZEN MEDIA**, collaborative research project (http://ist-citizenmedia.net)
- **ohrenstrand.net,** das Berliner „Klangstrandbad“ (http://www.ohrenstrand.net)
- **Inventionen** music festival, Berlin (http://www.inventionen.de/)
- **5th Sound and Music Computing Conference** (SMC08), Berlin (http://www.smc08.org/)

6.1.2 Crew Köln

- **Dipl.-Ing. Frank Melchior**, Fraunhofer Institut Digitale Medientechnologie IDMT: wave field synthesis consulting
- **Jörn Nettingsmeier**, freelance event engineer: executive producer in Köln, networking and stream supervision, rigging and rope access.
- **Tonmeister Hans Schlosser**, assistant professor at RSH: lead sound engineer and recording supervisor

Students of Sound and Video Engineering at RSH:

- **Simon Hildenbrand**: remote preamp and HD recording operator
- **Katrin Körber**: microphones, cabling, photo documentation
- **Jindrich Maßner**: microphones, cabling
- **Benjamin Simon:** microphones, cabling
- **Simon Spillner**: cabling, camera operator, sound truck driver
- **Josuel Theegarten**: microphones, cabling
- **Michal Zöllner**: sound truck operator

6.1.3 Crew Berlin

- **André Bartetzki**: Assistant WFS sound operator
- **Thilo Koch**, research associate: WFS cluster administrator, network engineer
- **Andreas Lehmann**: network support
- **Prof. Fernando Pablo Lopez-Lezcano**: streaming supervision
- **Eddie Mond**, research associate: WFS decoding supervisor
- **Wilm Thoben**, research associate: co-ordination
- **Tonmeister Prof. Dr. Stefan Weinzierl**: Project supervisor and WFS sound operator

6.2 Technical information

6.2.1 Dom network

The local network at the Dom consisted of a 100 Mbit/s Ethernet switch in the sound truck, whose uplink was connected to the laser on the spire via a media converter and 150 metres of fibre.
We used a 4-way fibre cable drum by ALVA Audio (marketed as a MADI cable), with two live and two dark fibres for backup. With its light weight and very rugged sheath, it was well suited to the job; after two weeks in the open, it showed no UV damage or strain marks from wind load and was easy to handle.
The installation was done from top to bottom with the drum slowly unwinding as it was abseiled to the ground. This mode of installation was chosen to reduce the strain of pulling the fibre over a significant length.
In places where the fibre had to be run over sharp edges, pieces of rubber hose cut lengthwise were slipped over it for extra protection. The cable was anchored every 2 to 5 metres with elastic straps.
Thanks to the optical connection, we did not have to worry about hum and signal degrada-

Illustration 14: cable drop from the intersection spire, with rubber edge protection.

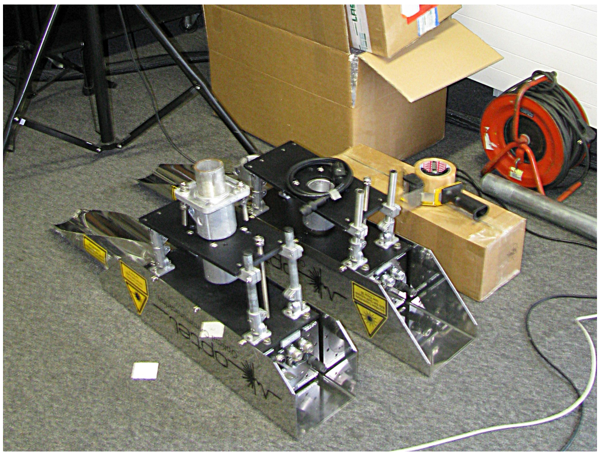

Illustration 15: The components of the laser link, ready to be shipped back after the show.

tion or take safety precautions against lightning.

The laser link comprised two Opticom FDD units (by OPTEL), with a raw bandwidth of 155 Mbit/s. It is fully transparent and behaves just like any other 100BaseFX segment. End-to-end ping times were on the order of 0.2 msec, and even at 95Mbit/s load, packet losses were well below 0.01% under normal weather conditions.
On the KHM side, the laser was patched through to a Linux NAT router which was in turn sitting on the CitizenMedia gigabit backbone.
After another 4 hops, the data entered the DFN X-WIN (german scientific backbone) and was finally transmitted to Berlin.

6.3 Software

Both the transmitting and receiving machines were Linux systems with fully pre-emptive kernels (using the current realtime patches). Sound routing was handled by JACK (http://jackaudio.org)

The original plan was to use JackTrip (http://ccrma.stanford.edu/groups/soundwire/s oftware/jacktrip/) at revision 147, since Prof. Weinzierl's institute had used it for previous smaller projects with very good results.

JackTrip allows for redundant transmission of audio packets. Unfortunately, this is accomplished by grouping several audio chunks into single UDP packets. For instance, at redundancy level 3, audio chunks would be transmitted in groups of three as [1,2,3], [2,3,4], [3,4,5], [5,6,7] etc., increasing packet size by a factor of three. Thus at 24 channels, the MTU limit will effectively prevent redundancy levels higher than 1.

We had briefly toyed with TCP tunneling to secure the transmission against packet loss, but decided not to follow that path, because listening tests had shown that the dropouts from occasional packet losses (very rare during our tests) were generally not audible or at least unobtrusive, and the TCP tunneling software we used was quite old and fragile.

To deal with out-of-order packets (a very frequent phenomenon when streaming UDP over public networks with redundant routes), we used very large input buffers to give the decoder time to rearrange the data before it was actually needed.

As described above, the network conditions on the rehearsal and show days prevented us from using JackTrip. The JackTrip author is currently working on TCP support that should help alleviate such problems. JackTrip development has been made public as of Aug 14 2008[2].

The fallback solution was implemented using the current Icecast stream server release[3] and a custom version of the IceS stream source maintained by Karl Heyes[4], which includes JACK support.

The decoding was done with MPlayer, its ALSA backend directly feeding an RME HDSP/Multiface. For some reason, we could not find a JACKified player that would correctly handle multichannel Vorbis streams and open the correct number of JACK ports. It appears that most implementations silently assume two-channel use.

In the near future, both JackTrip and recent netjack developments will be viable alternatives to the Vorbis-encoded streams we used, if sufficient bandwidth is available and their redundancy and buffering schemes perform as designed.

Multichannel Vorbis is a very useful tool today, and the specification appears to allow for up to 256 discrete channels. However, clients supporting multichannel streams are still somewhat scarce and clients should be closely evaluated before production use.

2 Project page is http://sourceforge.net/projects/jacktrip/.

3 Version 2.3.1, http://www.icecast.org)

4 http://svn.xiph.org/icecast/branches/kh/ices

6.4 References

6.4.1 Related publications

Hans Schlosser, Frank Melchior, Jörn Nettingsmeier, Stefan Weinzierl,"Mer losse d'r Dom en Kölle" - Weltpremiere einer Wellenfeldsynthese Live-Übertragung. Workshop at VDT International Convention, Leipzig 2008.
http://stackingdwarves.net/public_stuff/linux_audio/tmt08/Workshop_TMT08.pdf

6.4.2 Images

Illustrations 1,3-6, 9, 13-15 by the author; 11-12 by Frank Melchior; 2, 7 and 10 by Katrin Körber; 8 by Michael Zöllner, based on a floor plan provided by the Dombau-verwaltung.
Used with permission.