Adsonore

NATASHA BARRETT

Blåsbortvn 10, 0873 Oslo, Norway
E-mail: natashab@notam02.no

This paper explains the concept, technology and reality of the sound installation *Adsonore*. The important aspects of *Adsonore* are that it is interactive, continuously changing, permanent and located in a public space. The functioning of the human immune system lies behind its interactive design, the real-time sound processes and the physical construction. *Adsonore* is therefore directly related to the building in which it is housed – a new building for biological basic research.

1. INTRODUCTION

Adsonore is a permanent public space interactive sound installation located in the ten-floor stair tower of the new building for biological basic research (Bygg for Biologiske Basalfag) at Haukeland Hospital in Bergen. *Adsonore* was designed and built by the author with the help of Øyvind Hammer. Kjetil Matheussen at NOTAM provided valuable LINUX support. The commission was a bold initiative from the National Foundation for Art in Public Buildings, Norway in 2002, tested throughout 2003 and accepted as complete in 2004. As with all commissions from the National Foundation for Art in Public Buildings, the work requires a life of at least twenty years.

Adsonore is inspired by the human body's immune response. Instead of pathogens (bacteria, viruses and other foreign invaders to the human body), motion and sound created by visitors trigger *Adsonore*'s activity. This paper explains the installation, the technology used, and the passage from concept to reality. Interested readers can find a good introduction to the human immune system in Lodish *et al*. (2000).

The new building, designed by Narud-Stokke-Wiig, is located close to the top of one of the highest hills overlooking Bergen. The stair tower, with glass windows from ground to top floor, provides a panorama over the town (figure 1). When the National Foundation for Art in Public Buildings commissioned a work for this location it decided that a visual artwork would distract from the view and that a sound installation would be more appropriate.

The National Foundation for Art in Public Buildings began the commissioning process before the building was complete. Nevertheless, due to the duration of this process and the most difficult and lengthy task of gaining acceptance from the employees who would be using the new building, it was not possible to 'build in' a construction foundation for *Adsonore*. This meant that all cabling (close to 4 km) and loudspeaker equipment needed to be mounted on the public side of the walls and ceiling. Therefore, unable to make invisible the audio equipment, *Adsonore* took on an important visual element.

Designing and building a site-specific sound installation that is to run continuously in a public space over many years raises a number of aesthetic and practical questions. It is necessary to find a concept relevant and interesting to all those using the space on a level that is more than decorative, and then to find a way to fulfil these aims in practice. The work needs to provide an accessible and immediate impression for the first-time visitor, yet evolve in a way that will satisfy closer inspection by people using the space for days, weeks, months and years. Furthermore, there are a number of considerations demanded by the space: the location is used every day by employees, students, visitors and patients, whether or not they are interested in *Adsonore*. In everyday use, most people would use the stairs to cross one or two floors, while for greater distances they are more likely to use one of three lifts. Heavy fire doors are located at the exit from each floor, significantly isolating the sound within the tower, but a design flaw in their construction results in a heavy 'bang' when falling closed. Nevertheless, this isolation of sound to the tower area means that the choice of *Adsonore*'s volume is made in connection with the installation space and the people using it, rather than in consideration of adjacent spaces.

Using scientific principles as the foundation for an aesthetic abstraction can be loaded with problems. For example, if one uses cutting-edge science, there is the danger of current principles changing with further research, while if one uses accepted and established subject matter, in biology one is normally dealing with relatively straightforward principles, and after a few months an aesthetic abstraction may be seen as boring or mundane. Despite these problems, I chose an area of basic biological study upon which to find abstractions for *Adsonore*. I use the term 'abstraction' because *Adsonore* is not a model of the immune system, but rather takes principles from the body's immune response as a rule-base for the fundamentally

Organised Sound **10**(2): 111–119 © 2005 Cambridge University Press. Printed in the United Kingdom. doi:10.1017/S1355771805000750

https://doi.org/10.1017/S1355771805000750 Published online by Cambridge University Press

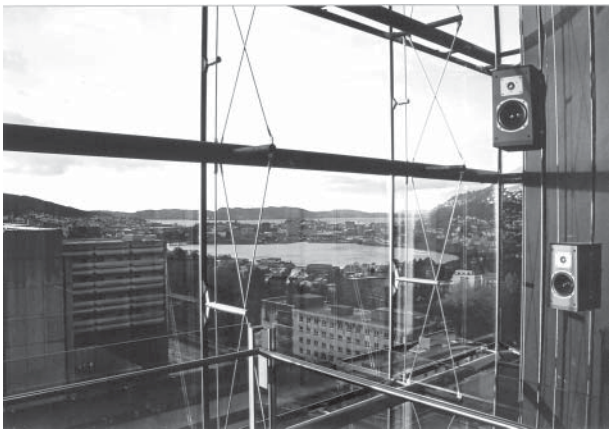


Figure 1. View from upper floors of the west stair tower (photo: Sekkengstad 2004).

interactive nature of the project. As science does not completely understand the mysteries of the immune system, *Adsonore* draws on established facts, and any diversion is through aesthetic development rather than changes in modern biological discovery.

Sound examples accompanying this text are taken from live activity on 10 and 12 November 2004.

2. INITIAL TESTING

Most stair towers are characterised by a boomy, highly coloured acoustic. The location for *Adsonore* was no exception. The budget did not allow for any acoustic treatment, and so *Adsonore* needed to take advantage of, rather than suffer from these acoustic oddities. Acoustic tests to find the frequency ranges and physical locations for the most prominent resonant areas were undertaken. The results of these tests heavily influenced the choice and location of loudspeakers and microphones, and many of *Adsonore*'s real-time sound processing parameters. Loudspeaker types and their locations were selected to combine directionality and acoustic clarity – allowing clear spatialisation – with carefully controlled resonant 'blooms'. These blooms were enough to allow the space to 'sound' while avoiding acoustic feedback. The microphones were chosen and located to record a broadband and reasonably dry signal while being out of reach from accidental physical contact.

3. VISUAL CONSTRUCTION

When you enter *Adsonore* you see a 'chain' of fourteen loudspeakers (Quadral Argent 21) hanging in staggered pairs along the vertical area of the tower (figure 2). These loudspeakers provide the means for sound to travel rapidly or 'hover' at each floor level. Each pair of loudspeakers is positioned to provide a localised area of sound and also to spill sound at lower volumes to adjacent floors. As you walk up or down the stair you will see small loudspeakers (10 watt

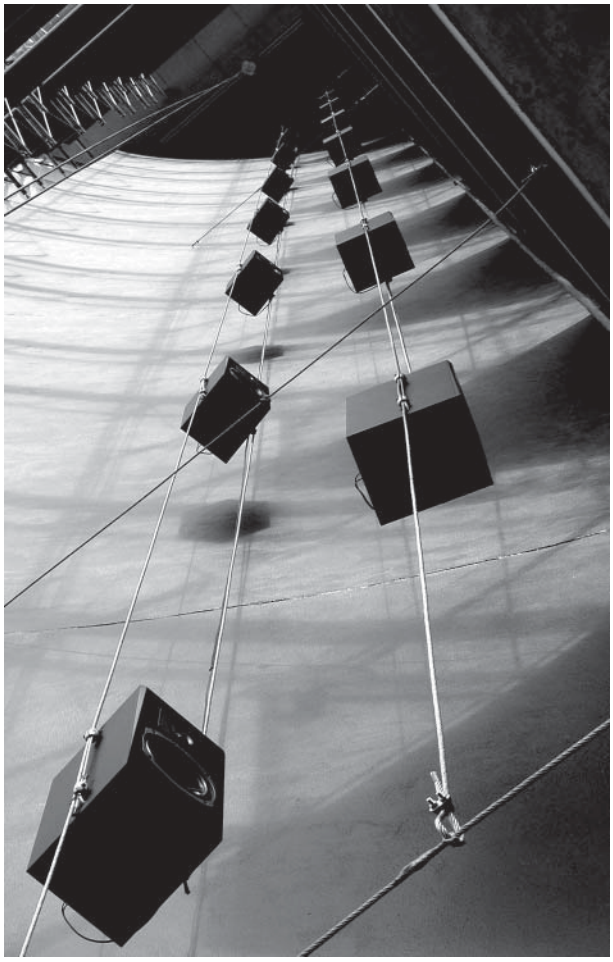


Figure 2. Main loudspeakers.

broadband speakers) mounted in triangular fixings above head height, of which there are forty-two in total (figure 3). These provide localised sound and allow material to be spatialised along the stair. On each floor, at waist height, are located three infrared sensors. One sensor (Omron E3Z) is located by each door and detects the entrance or exit of people to the tower. When you continue up or down the stairs, on each floor you will see a grey box and two reflectors. The box contains two infrared sensors (Omron E3F2) pointing in opposite directions to detect motion information in the stair area. A boundary layer microphone (AKG C 547 BL) is mounted above head height in the area of each pair of sensors (figure 4). Although initially considering hidden microphones, the invasion of people's privacy could be problematic even if one notifies the visitor.

For a visually acceptable and safe solution to the cabling problem, steel tubing was mounted along the ceiling edge of each wall. Accepting that it was impossible to completely hide the audio cables, these tubes tidy the cables, but allow them to be seen – as the veins and arteries of *Adsonore*. This visual and practical problem was solved in collaboration with the architects.

https://doi.org/10.1017/S1355771805000750 Published online by Cambridge University Press



Figure 3. Small loudspeakers.

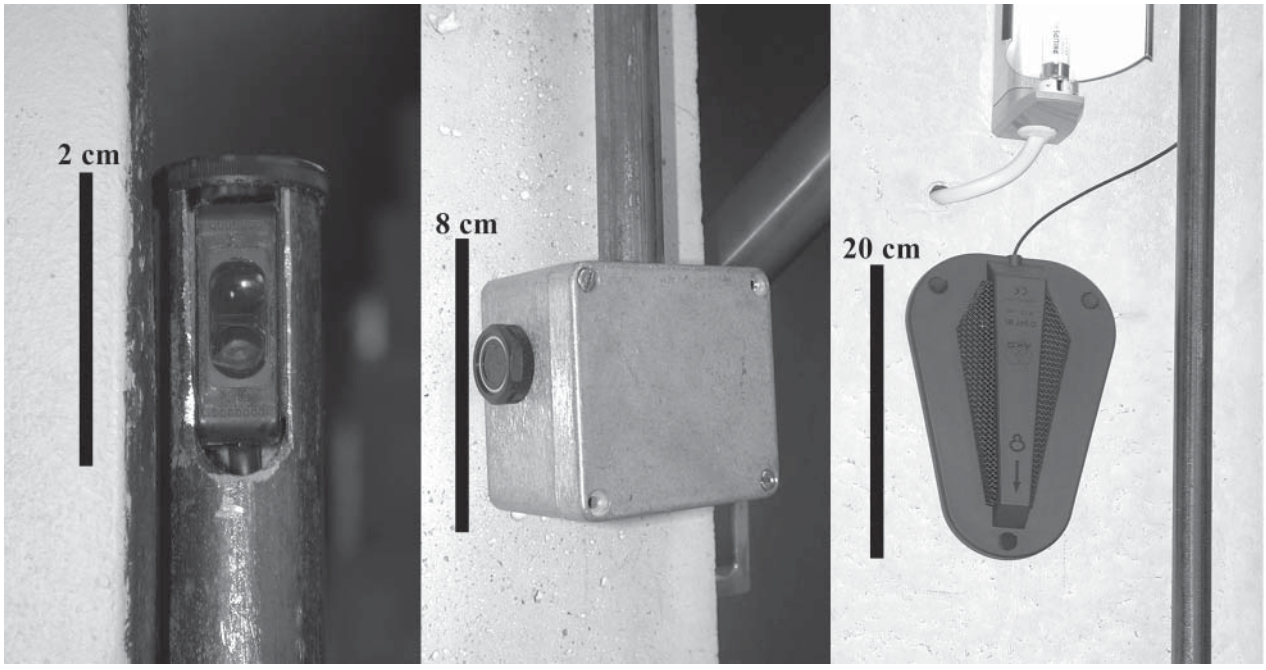


Figure 4. Sensors and microphones.

The information from microphones and sensors is analysed by the computer, after which different types of processing and spatialisation are applied to sound played back over the fifty-six loudspeakers. Depending on the state of the system, people entering the tower may be listeners to responses already activated, or may themselves activate a response. Every person activating a response in *Adsonore* leaves a trace influencing its future development.

4. IMMEDIATE SOUND IMPRESSION ON THE LISTENER

It is not necessary for the first-time visitor to have any understanding of *Adsonore*'s operation, or of the human immune system. The only requirement is an appreciation of sound rather than expecting the traditional instrumental concept of 'music'. *Adsonore* 'playing' with the visitors' own sound is enough to

https://doi.org/10.1017/S1355771805000750 Published online by Cambridge University Press

tempt those interested into deeper engagement. It does, however, help to know three pieces of information, which are provided in small signs mounted at various floor levels:

(1) To interact you need to make sound at the correct time and of 'intentional volume'. In reality this means a volume higher than the sound of normal walking or of quiet speech. This is to allow those not interested to pass silently through without activating a response. It also allows interested people to 'play with' or 'trick' *Adsonore*.
(2) Most of the sound materials are recorded from sound made by visitors. Therefore the variety and activity of *Adsonore* is, over the short term, up to the individual, and over the long term up to visitors as a collective.
(3) If many people occupy the same location of the tower, interaction becomes a group activity. This means that at different times of day a visitor will gain a different impression of how the system works. The clearest impression is gained in quiet periods when only one or two people are in the tower. At this time it is clear what response each individual is activating. During busy times – such as early morning, noon and late afternoon – the quantity of people moving through the tower activates a complex layer of simultaneous processes. These times of day can be fun and exciting, but only regular visitors will begin to understand what is happening.

In the 'beginning', which is the point *Adsonore* is initially switched on or reset, *Adsonore* encounters all sound for the first time. Gradually it 'learns' and responds more quickly and in new ways.

5. TWO TYPES OF INTERACTION

Interactivity involves a two-way passage of information. People moving to a goal will trigger motion sensors, and if the sound they produce is above a certain volume, their sound will be recorded, stored, processed in real time and played back. Visitors may unintentionally produce sound at a volume to trigger recording, and the reaction of *Adsonore* may influence their behaviour. We can call this unintentional interactivity. In other words, visitors unintentionally provide input to the installation, and the subsequent reaction of the installation may affect their behaviour. On the other hand, a person may intentionally move or produce sound in an attempt to begin an exchange with *Adsonore*. This can be called intentional interactivity. In such situations, people can try to fool the system by making unusual sound, by changing their physical motion speed or direction, or by tricking the system into attacking itself ('autoimmune' response). In a space used every day for purposes other than interactive art it would be interesting to think that both intentional and unintentional interactivity are important elements, where unintentional interactivity can stimulate the visitor into intentional interactivity on the same or future visit. I have experienced this link between intentional and unintentional interactivity as being common in many public-space sound installations. However, with *Adsonore* I found that unintentional interactivity stimulated a negative response in the visitor and it was necessary to dampen *Adsonore*'s reaction sensitivity.

6. THE INTERACTIVE PROCESS AND THE MODEL

The immune response consists of a complex set of mechanisms, involving different types of cells, antibodies, tagging, feedback and memory processes, all working together to identify and destroy a foreign item. These different mechanisms occur simultaneously. Likewise, the different functions of *Adsonore*'s system occur simultaneously. However, for purposes of description, these functions are below divided into stages, each stage derived from its position in the hierarchy of the controlling program.

Although there are many different types of cell working in the immune system and five types of antibody protein, there are three cellular processes and one antibody process relevant to *Adsonore*: the disabling of foreign bodies, the production of antibodies, the stimulation of the reproduction of different types of immune response cell, and the 'tagging' of foreign bodies to be identified and disabled. In *Adsonore* each of these has sound and motion analogies.

The processes listed below can occur simultaneously and independently on seven separate floors of the tower, as well as the sounding results being spatialised across floors. All floors are united through contributing to, and accessing, a common data bank of information and sound. Each process involves variables that were calibrated during many months of testing. In this way, the controlling program, and inevitably the composition, could be completed in the studio and then tuned to function in the reality of the installation space.

6.1. Sampling

(1) Sound from each microphone is filtered to remove some of the constant background noise found in the tower. Technical equipment from research facilities, as well as ventilation systems, creates a continuous hum that is partly airborne and partly structurally carried.
(2) An envelope follower detects if the sound input rises above a threshold. Only when held above this threshold for more than 100 ms will sound

https://doi.org/10.1017/S1355771805000750 Published online by Cambridge University Press

be sampled. The hold duration is to prevent transients such as heavy footsteps unintentionally triggering the record function. The acoustic provides the necessary sustain for rapid articulations such as clapping or foot tapping to trigger recording.

(3) On each floor are located three infrared retroreflective sensors, as described above. The two sensors mounted across the stairs provide motion direction information and activate the record function. When a sensor beam is broken it tells the program to record sound if the volume threshold has been exceeded. The duration of the recorded sound is between 2,000–7,000 ms. If motion direction is clear (see below), the 'on-off' information determines the duration of the sound bite between these values. If motion information is ambiguous, the maximum 7,000 ms duration is used. This maximum value has been set for three reasons: (i) to allow a large number of sounds to be active at any one time, (ii) to act as a default time value if the on/off data from the motion sensors are ambiguous, and (iii) 7,000 ms is a long enough duration for some interesting sound bite, but short enough to avoid lengthy extracts of possibly private conversation being recorded.

(4) The recorded sound is analysed by the 'audio classifier'. The audio classifier is explained in section 9. If classified as a new sound, it is normalised and saved to disk with a unique file name and used in the next stage of processing. If classified as an old sound, *Adsonore* already has 'antibodies' to this foreign body. In practice this means that the computer retrieves the old sound from its catalogue and processes it in a way analogous to the 'lock-in-key' interaction between antibody and antigen (see processing layer 2). Although normalisation reduces interesting amplitude differences between the source recordings, it is a necessary stage to gain control over the output volume of different layers of processing. Normally these amplitude differences can be used advantageously, but in a public space the acceptable dynamic range is much less than that of a concert or temporary installation.

Sound example 1a is a commonly occurring sound, 1b a less common sound. Examples are in mono with background information.

6.2. Sound processing and spatialisation layer 1

From the sensors crossing each floor, the computer calculates the order of 'on-off' messages and their delay times. For example, if the triggering of sensor-a follows sensor-b with an interval between 2,000 and 4,000 ms, we can with reasonably certainty calculate the direction of motion of one person. 2,000–4,000 ms is a reasonable duration within which one person would cross the sensor boundaries. If, on the other hand, the delay is less than 2,000 ms or greater than 4,000 ms, or even if one sensor is triggered twice consecutively, motion information is ambiguous. In this instance we can assume that a visitor is trying to fool the system, or that there are many people in the tower moving in different directions. Even though the sensors are located at waist height, it is possible for moving arms to create fast, double activations. Therefore repetitions less than 100 ms apart are ignored.

Recorded sound is delayed by 1,000 ms and played in its raw format over the closest small loudspeakers. The delay is to prevent feedback occurring through the microphone and the two adjacent small loudspeakers. If clear motion information is received from the sensors, the sound is spatialised in the direction of this motion, for example along small loudspeakers 1-2-3 or 4-5-6 in figure 5. The speed of this motion is determined by the delay between sensor activations. If motion information is ambiguous, the sound remains stationary on small loudspeakers 3 and 4 in figure 5. This first reaction is *Adsonore* recognising a foreign body and trying to dispel it from its system, just as the immune response involves antibodies that move in random but continuous trajectories 'tagging' the foreign body.

If many people are entering the tower, occasionally the recorded sound will be processed with a resonant filter effect, producing a ringing harmonic quality, as if the immune system is struggling under the strain of its invaders.

6.3. Sound processing layer 2

After 2–7 seconds, sound reaches the landing. Here it is looped and played for 10 seconds over the two main speakers for that floor in a left-right oscillating motion, at a speed randomly variable between 0.1 and 100 Hz. The purpose of this motion is to allow the visitor to clearly perceive the sound above any background noise and acoustic bloom, without the need for high volume. If the sound is classified as an old sound, the original is retrieved, processed with a harmonisation effect and played over the main loudspeakers instead of the sound recorded by the microphone. To reduce the degree of transformation in this early stage, randomly variable parameters of the harmonisation (transposition, delay, window size) are restrained within a narrow range. Here we hear one side of the sounding analogy to an antibody: a 'negative', in some sense, to a fragment of the foreign body sound – similar to the 'lock-in-key' interaction between antibody and antigen. Sound example 2a is an extract from this layer of response to a 'new' sound; 2b is an

https://doi.org/10.1017/S1355771805000750 Published online by Cambridge University Press

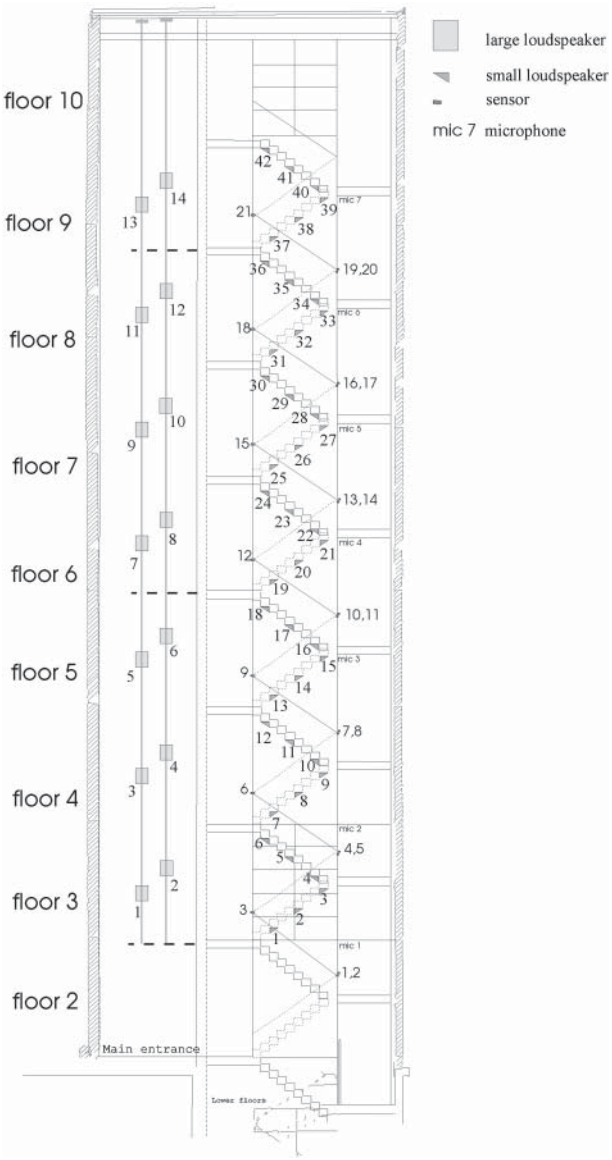


Figure 5. Overview of installation visible equipment.

example of a response to an old sound. Both examples are in stereo and recorded direct to disk without ambient sound.

6.4. Memory log

If the recorded sound is identified as 'new', its identification number is stored in a temporary log file unique to each floor. If the sound is not heard again within twenty-four hours it is cleared from this log. This log clearance is related to the antibody analogy in section 6.5 below – where our body breaks down antibodies if they do not encounter an antigen within a few days. If within this time span *Adsonore* detects the sound a second time, it is classed as 'old' and stored in a permanent log file common for all floors. Every twenty-one days this list of sounds, or what can be called the 'memory file', is stored to disk and its contents retrieved for sound processing layer 4. The contents of these memory files are analogous to the body's own memory cells, which result in our immunity to different infections and diseases. Sounds from temporary and permanent log files are accessed by sound processing layer 3 (see below).

6.5. Sound processing layer 3

The processes responsible for antibody variation in the human body are important. A two-layer random process achieves infinite variation: random variation of genetic source segments and random transcriptional splicing of these segments. A similar two-layer random process of transformation and splicing transforms the sound-antibody.

After a delay time of 15 seconds, which is sufficient to join onto the end of first- and second-layer processing, old and new sounds are retrieved from the log files for use in the 'splicer'. Each floor has its own 'splicer' to handle floor-specific sounds. The splicer is a granulation program designed to imitate the basic antibody RNA splicing process, where small segments of sound are cut from randomly chosen new and old sound files. Some of the variables include pitch transposition of segments and time variation of segments as a percentage of the source sound duration. If *Adsonore* detects 'old' sound activity, the size of the splice window reduces over a five-day period resulting in increased tension to this sound layer. This period is the duration over which our own immune response reaches a peak in its defence against an antigen. The system then subsides over a sixteen-day period. The splicer output is heard for 15 seconds on the left-hand main speaker. The seven left-hand speakers can therefore be regarded as 'b-cells' or antibody factories. Sound example 3a is an extract from this layer when in a slow phase; 3b is an example one day later.

6.6. Spatialisation 2

Concentrations of immune cells are found in the lymph system and guard the portals of entry into the body. Antibody and 'killer-cell' activity is focused in these areas. In *Adsonore*, the entrance doors on each floor are the points at which people enter. When sensors at each entry door are triggered, the output of all splicers or 'b-cells' is spatialised vertically through the tower in the direction of the trigger points. This is a rapid process lasting only 900 ms. Sound example 4 is an extract where all splicer outputs form one fast moving body; extract from peak in splicing cycle.

6.7. Sound processing and spatialisation layer 4

When an old sound is detected, it is immediately retrieved and processed with some extreme effects:

https://doi.org/10.1017/S1355771805000750 Published online by Cambridge University Press

wider parameter harmonisation than earlier, amplitude modulation, spectral inversion and convolution. The processed sound can be loud, 'full', and move rapidly over small and large loudspeakers in the direction of any motion information. If the motion information is ambiguous, this is one instance when the sound is spatialised in the direction of the last clear detection. Processing and spatialisation in layer 4 happens as soon as the computer identifies the nature of the sound. This means the reaction is almost immediate – happening before or simultaneously with sound processing layer 1. Contrary to its design and function, this layer of processing has been significantly reduced in volume due to complaints by some employees working in the building. Some found *Adsonore*'s response disturbing; others were concerned that it could be perceived as frightening. Sound examples 5a, 5b and 5c are extracts from this layer.

6.8. Sound processing and spatialisation layer 5

Every twenty-one days a log file of 'old' sound names called the 'memory log' is written to disk. These log files accumulate infinitely. A new file is randomly retrieved on a daily basis from the available pool of 'memory files' and the contents used to recall 'old' sounds. Processing layer 5 is a parallel to the 'memory cells' of the immune system. When the door sensor is triggered, a sound is retrieved and processed with longer, continuous developments than in the previous stages. Each sound is continuously looped and processed with a randomly controlled choice of reverberation, delays and high-Q resonant filters of numerous different frequency ranges. This processing stage produces a very quiet result heard just above the natural background noise level, starting at the floor on which the door sensor was activated and moving slowly up and down the main loudspeaker chain. One sound may be looped for up to four minutes, and thus at most times of day *Adsonore* will constantly 'murmur'. The looping and transformation effects take advantage of the tower's natural acoustic to produce a pulsating drone. Some people have suggested this makes *Adsonore* 'alive'. Sound example 6 is an extract from this layer.

6.9. Pre-prepared sound retrieval and spatialisation

Even if the visitor presents a wide variety of sounds, there are restrictions placed on the resulting sound world. With our current technology there are some limitations conditional on the speed of the computer processor and the speed of the hard disk. However, there are three more important issues independent of any technological limitations:

(1) Although the natural sound and acoustic are characteristics of the space that *Adsonore* has been designed to accommodate, their presence in all live sound bites may lead to long-term problems of sound variation.
(2) What can be achieved 'out of time' in the studio will inevitably be quite different from the creations of non-intervention real-time processing. It is this control over spectrum and morphology that is lacking in any 'stand-alone' real-time processing I have been able to create so far – even in live performance contexts.
(3) To be faithful to the concept, the only sounds *Adsonore* should work with are those made by visitors.

In reality, the imagination or willingness of many regular visitors can be limited. Paradoxically it is these very people that are most quick to complain when they perceive something as lacking enough variation. For *Adsonore* to function, it is necessary to at least satisfy, if not enrich the lives of those uninterested. For these three reasons it was necessary for *Adsonore* to be able to access pre-composed sound material. In total, 980 prepared stereo sounds of duration 0.5–120 seconds were created. Male and female vocal materials, as well as some internal body sounds such as rumbling tummies, were used as source materials and in turn crafted to create a variety of attacks, textures, and developing morphologies. A separate audio classifier is 'primed' with analysis data from these pre-composed sounds (see section 9 below). The classifier compares incoming live sound to the 'primed' log file, which selects pre-composed sound most closely matching the incoming information. The probability that a pre-composed sound will be retrieved is determined by how far the system is through its twenty-one-day immune response cycle. If the cycle is at its five-day peak, then there is a greater probability that a pre-composed sound will be selected for playback. This probability decreases again over the remaining sixteen days of the cycle.

Pre-composed sounds are played on the main speakers at the floor on which the live sound was detected. In addition, a local single small speaker may be turned on until another process turns it off. Sound examples 7a, 7b and 7c are examples of prepared sound material.

7. THE TOTALITY: COUNTERPOINT, PUNCTUATION AND SIMULTANEITY

In a period of maximum activity, *Adsonore* may be responding to simultaneous stimulus from all floors. The acoustic separation between floors is enough for a listener to hear their personal effect in the foreground, while hearing the effect of others resonating in more distant areas of the tower. As the activity reduces, a clearer counterpoint is heard between different floors.

https://doi.org/10.1017/S1355771805000750 Published online by Cambridge University Press

In the background, the pulsating loops of the 'memory sounds' may be audible. When only one or two people remain, the slowly moving memory sounds form a background wash, while any real-time interaction creates punctuating and clearly audible stages in the *Adsonore*'s response. Sound example 8 is an extract from 10 a.m. where simultaneous activity from three to four different floors are mixed to mono for the purposes of this text (sound example recorded direct to disk).

8. DURATION

Adsonore should run continuously day and night. However, some people, surprisingly to me, said they were scared to use the stair tower at night because of the sound. Therefore, *Adsonore* now only runs during office hours.

Adsonore will temporarily shut itself off if it detects continuous high volume over four or more of the seven microphones. The original purpose of this program is in the event of a fire alarm, but also serves as an automatic 'safety switch' if any of the audio hardware were to fail.

Adsonore can be reset anytime by running a simple script. This is at the discretion of the university computer department.

9. TECHNOLOGY

When asked to build a computer sound installation that could run for twenty years, the immediate questions concern software. Hardware will inevitably fail and most likely become obsolete within this duration, and so it was necessary to choose an operating system and sound program that could be recompiled for the next general hardware. It was also necessary to use software reliable enough to run for months and even years without 'crashing', and a system that would allow me to sit comfortably in Oslo during the testing period – or in fact anywhere in the world – visually controlling the *Adsonore* machine located in Bergen over the Internet. LINUX and Pure Data were the natural choice. In the short term the choice of LINUX caused an array of frustrating problems in connection with unusable sound card drivers. I am indebted to Kjetil Matheussen and Gunter Geiger for their help with the Hammerfall HDSP9652 ALSA driver. In the long term, I hope this choice of software will allow a complete overhaul of the hardware when the need arises. There is, however, no long-term maintenance budget, and the daily operation of *Adsonore* is in the hands of the University in Bergen computer department.

For additional durability, the computer uses a hot-swap dual power supply and runs from a UPS power and surge protection unit.

9.1. The audio classifier

One of the most technically challenging components in *Adsonore* is the audio classifier. This piece of software receives sound in real time, compares it with sounds in a library, and either recognises the incoming sound as having been encountered before, or as a new sound. The audio classifier is thus a central part of the 'immune response' of *Adsonore*.

Perfect audio classification requires a highly complex algorithm. Because this is still an active area of research in digital signal processing and artificial intelligence, we designed a somewhat simple system that was realistic to implement in real time. The audio classifier is programmed in C and compiled to run under Pd within the main software. It involves a three-stage process:

(1) 'Essence' – a real-time FFT (phase vocoder) implementation analysing incoming sound and outputting real-time amplitude, pitch and brightness (spectral centroid) information.
(2) 'Vecstat' – analyses the multivariate output stream from 'Essence' and calculates nine parameters over a time chunk: amplitude mean, amplitude variance, amplitude dynamics (mean of first-order differences), pitch mean, pitch variance, pitch dynamics, brightness mean, brightness variance, brightness dynamics. As it is necessary to reduce the dimensionality of the problem by parameterising the sound, tests signals were analysed with respect to a larger number of parameters, and principal components analysis used to decide on the most relevant of these.
(3) 'Statlib' – a library of statistical parameter vectors that stores information from 'Vecstat'. 'Statlib' uses so-called unsupervised learning, where the program continuously updates itself according to new incoming sounds. When the distance between an incoming parameter vector and its closest match in the library is greater than a threshold value, 'Statlib' stores the incoming vector, sends out a 'new' message and a log number for that vector. The threshold value is user controllable so that the sensitivity of the response can be tuned, and is given in units of proportion of variance of vectors in the library, which is dynamically changing. When the distance between an incoming parameter vector and its closest match is less than the threshold value, 'Statlib' sends out the message 'old' and the matching log number. The new log number is used as a name for new stored sound. The old log number is used to retrieve an existing sound.

'Essence', 'Vecstat' and 'Statlib' were written by Øyvind Hammer.

Two pieces of hardware were custom built: (i) a 4-in-42-out computer-controlled audio matrix distributor

https://doi.org/10.1017/S1355771805000750 Published online by Cambridge University Press

to spatialise sound over the small loudspeakers, (ii) voltage to MIDI converter for the twenty-one infrared sensors. Although it was possible to buy the computer-controlled audio matrix distributor, its cost and size were enormous. Design and construction of our own device meant that it was less expensive, did exactly what was needed, and most importantly, in the unlikely event of its failure, could be fixed by anyone with a little electronics knowledge. The same was true of the voltage to MIDI converter, although this box already existed in the form of the NOTAM MIDI box. Øyvind Hammer designed and built both of these devices. All computer hardware is housed in its own room on the top floor of the tower.

10. CONCLUSION

Adsonore is a large installation. Taking it from concept to reality was a challenge. The need to handle budgets, select and buy every screw and cable, and hire appropriate labour is a task I would not rush into alone a second time. The demands of a public space such as stability, security, variability, and reducing the response of *Adsonore* to uninterested passers by have been satisfied without significant compromise to the concept. Nevertheless, even before the building was opened, *Adsonore* was a controversial project, opposed by many of the staff. After the work was opened, the 'user co-ordinator' for the building sent a questionnaire asking staff to rate the installation from 1 (bad) to 6 (good). Less than fifty per cent returned the questionnaire and surprisingly no students were asked to take part in the survey. From 127 replies, 95 gave a rating of '1'. Some of the added comments were negative and critical, staff being frightened or irritated by the sound, while others wrote 'cool', 'special' and 'out of the ordinary' (source of information from Sekkingstad 2004). Sekkingstad proposes several reasons for this reaction. He writes that the physical and confrontational aspect of the sound is problematic for some people. Others are annoyed that it was difficult to have a normal conversation in the tower – not because of the volume being too high, but because fragments of their conversation 'were suddenly thrown around the space'. Others were unhappy that the abstract electronic sound had no melodic or harmonic progression. Many people working in the building do not know anything about the human immune system and therefore find difficulty understanding how *Adsonore* functions. Sekkingstad further suggests that an underlying reason to the opposition may be one of social orientation. In Norway, sound has only recently found a place in contemporary art and society. People encountering sound art have not grown up with it as part of culture and society, have never visited a contemporary art exhibition and have no personal interest in that direction. The situation is somewhat different in other countries of Europe.

During the working period, the enthusiasm, support and co-operation from the architects who designed the building and also from the commissioning body were important. *Adsonore* has now been running for over a year and the general opinion, with some exceptions, is one of interest and enthusiasm by younger and temporary visitors, and a negative feeling of resignation amongst many members of departmental staff.

REFERENCES

Lodish, Berk, Zipursky, Matsudaira, Baltimore and Darnell. 2000. *Molecular Cell Biology*. W. H. Freeman.

Sekkengstad, Sr. 2004. *Kunstens Nye Stemme – noen betraktninger om lyd i samtidskunsten*. University of Bergen.

https://doi.org/10.1017/S1355771805000750 Published online by Cambridge University Press