THE SOUND GATE, A NEW ACOUSTIC CONCEPT

Alan D. Wallis, Olaf Alberts

Cirrus Research plc
Acoustic House
Hunmanby
YO14 0PH
United Kingdom

1. INTRODUCTION

The essential difference between a conventional Environmental Noise Monitor and a dedicated Airport noise monitor is that an airport noise monitor should distinguish aircraft from all other noise events and provide a unique measure of each aircraft event, with all other noise sources separated. What metric is used to describe the aircraft noise can be changed after data acquisition, as the computer in a current technology installation can convert one metric to another with ease, providing the raw data acquired is of sufficient resolution. Most current technology airport noise monitors use the technique of 'Short L_{eq}' as the acquisition medium, usually based on a 62.5 mS long basic integration period. This allows the computer to re-constitute both 'S' (Slow) and 'F' (Fast) responses to the stated tolerances of IEC 651 Type 0 and these can then be used to produce the various statistical values. Short L_{eq} also allows the system to recognise an individual flight by its time history and Figure 1 shows a typical trace at a local airport. Each individual aircraft is clearly defined and with such a trace, there should be no problem in identifying each one as an aircraft.

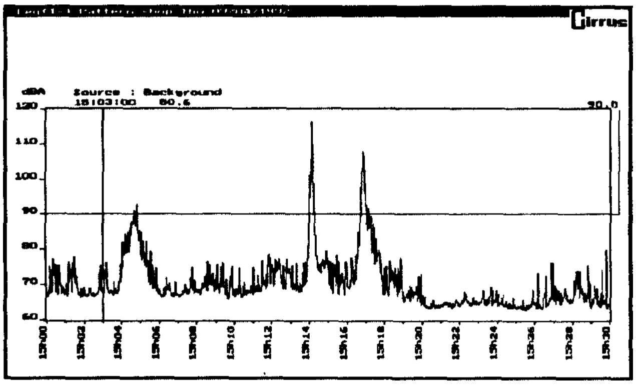

Figure 1: A typical time history trace

Of course, all monitoring points do not have such a good 'signal to noise' ratio as shown here and in a practical situation, each individual noise event may well be lost among other local noise sources such as vehicles or industrial noise, or even fail to emerge out of the background. One measure of an airport noise monitoring terminal, is how well it recovers the wanted signals in difficult conditions.

2. EVENT RECOGNITION ALGORITHMS

The algorithms used in the latest generation of Cirrus Research monitors, to recognise aircraft, have been described elsewhere (*ref 1*) and basically consist of a multiple threshold with specified event durations, rise and fall times as well as 'dwell' times; Figure 2 showing the general concept.

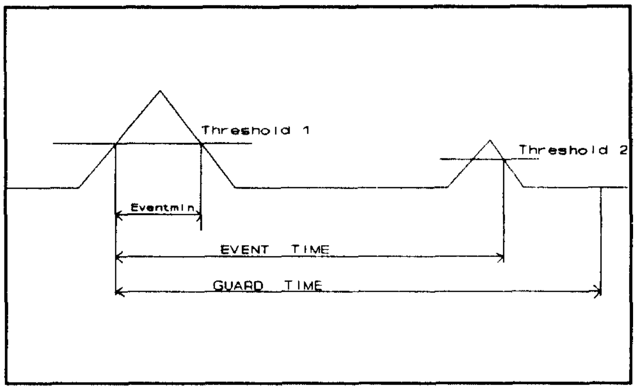

Figure 2: Event Recognition Template

Experience has proved that this complex multi-parameter algorithm will recognise about 99% of scheduled or military aircraft correctly if the site has a reasonable background noise. Like all other systems, it does have a lower success rate for very quiet aircraft or in areas of high background noise. The precise 'hit rate' or recognition accuracy of commercial systems, is a closely guarded figure for most companies, usually because they are so bad. For some systems in use today, where a simple threshold is used, even an 80% hit rate cannot be achieved over a long period and many flights are simply missed or the system reports a flight where none took place. Such low hit rates are a significant problem. For example, with 100 flights a day, 20 flights will be incorrectly identified and such a poor system cannot possibly be used as the starting point to penalise noisy airlines. Even with a 99% hit rate, there will be one flight per day incorrectly identified at such an airport and thus before any method of allocating noise to each airline can be implemented, something else must be brought in to reduce this recognition error rate, as at a major airport with over 1000 movements per day, even a 99% hit rate is totally unacceptable.

Some users have suggested building an FFT into the Noise monitor, using the spectrum as an aid to event recognition. The current Cirrus Research unit does have an FFT option, but in fact its contribution to the recognition process is very limited and certainly such an add-on has very limited value for this task. The main use of such a unit is to generate EPNL and similar indices. In a high accuracy system, the noise recognised at one monitor will be checked by a second or even a third monitor, so as to ensure that the noise really is an aircraft and not some local noise. As the monitors may well be separated by several kilometres only aircraft are likely to be 'heard' by more than one monitor. A noise source at ground level will probably not carry so far. This correlation process requires that all the monitors are connected 'on line' to the central host computer as it is here that the correlation between noise monitors takes place.

Recognising that a flight has taken place, at a particular noise monitor, is only the start of the problem, now it must be identified by carrier, usually using the flight number. Without this, you cannot set in place a method of penalising noisy operators. In other words, once a flight has been recognised it must be identified. What is needed is a method of tying an individual aircraft noise event to a specific flight and some current methods of achieving this are discussed.

3. BASIC APPROACH

Each airport has access to flight information, as used in public terminal displays and this is often utilised as one of the correlation parameters in identification. If the airport has significant General Aviation or private traffic, this may well be incomplete and significant flights such as a jet test flight will be missed. Many airports have a radar display and the data from this can be used as a system input, although in many places, this is forbidden. A derivative of this is to passively listen to the radar transponder fitted in the aircraft, but this suffers badly from false echoes in some installations. Yet a further method, although not in 'real time', is the use of control tower flight strips which are universally available and can be hand entered into the system. One thing that is reasonably well defined at an airport is the invoicing system, as without this landing fees cannot be correctly assigned and as today this is usually computerised, this can be used as a data stream. Many other methods have been tried from video cameras reading the aircraft tail numbers, to a special 'human noise monitor', who simply noted each movement and timed them accurately. In the real world these systems can identify that a flight took place, but none of them on their own can give the accurate identification of a particular noise event. Indeed, it is probable that some airports have penalised airlines where the wrong carrier was identified, particularly where radar was used to identify the flight.

Knowing that an event has occurred and that there is a flight at a similar time is still not adequate. The flight information data must now be tied to a particular noise event with no significant possibility of error. The problem with all these methods of identifying the flight is that they are all vague as to the exact timing. The exact time of take-off is not known and it is not a priority matter for the airport to know this. The time at which an aircraft pushes back from the stand is usually known as is the time it arrives, but many times a pilot will hold at the runway threshold for a minute or so and put the scheduled times in error.

The process used in RASP, the Cirrus Research Regional Airport System Program, is to correlate a noise event recognised at a particular noise monitor, with one of the many data inputs described above, depending on availability. This is done by allowing a known time to elapse after a takeoff, or before landing, to predict when the noise is expected at a particular noise monitor. This gives a 'confidence factor' depending on how far away from the expected time the noise is heard and the ratio of signal to noise at each monitor. This elapsed time can be reasonably accurate, as the take off and landing speed of jet transport aircraft is much the same no matter what type of aircraft is involved and for the much smaller, slower and usually quieter light aviation traffic, the correlation is, in any event, less important.

4. REDUCING THE ERRORS

With identification from one of the airport systems and a clear aircraft event from an NMT, there are still major problems. The main one being to get a true landing or take-off time to plug into the correlation algorithm. For example, if three flights go out together, the timing given by the airport system may well be in error by more than the time between flights. If this happens, the system may get the flight events 'out of step' and the error will not be recognised until a long period of quiet. In this case, a very low correlation confidence level will be generated which reduces the integrity of the system. To overcome this, Cirrus Research have developed and patented *(ref 2)* new techniques to accurately time the landings and take-offs using a 'Sound Gate'. The outline is shown in Figure 3.

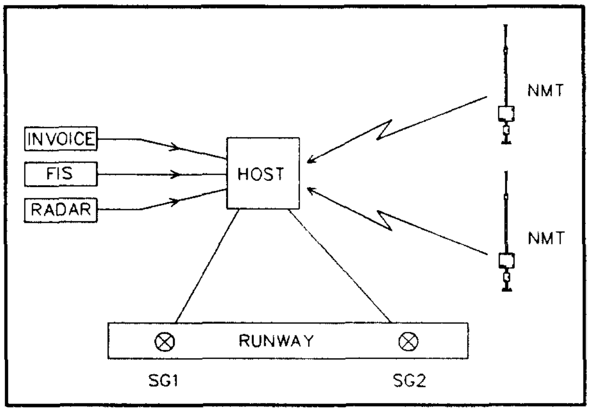

Figure 3: Complete system

The technique uses miniature noise monitors mounted into the runway centre-line near each threshold, with special filters which increase the signal to noise ratio of aircraft noise and select out some different engine noises. These miniature instruments are called 'Sound Gates' (pat applied for). Complex thresholds similar to those in the NMT, are fitted in the Sound Gates to recognise an aircraft event which can be timed to the nearest 62,5 millisecond.

The data is fed 'on line' to a host computer which takes the timed events from the Sound Gate, usually a CRL 247 and this determines:-

> The direction of travel, in other words the runway in use.
> Whether it is a take-off or a landing.

It also puts the aircraft into a crude engine size category by the maximum level and other parameters. As an example, on the centre line of the runway, noise levels above 130dB are unlikely to be anything other than a pure jet, but levels from 100 to 125 dB can be jet-prop or even a piston engine type.

Where the runway cannot be dug up to install the Sound Gate, a noise intensity probe unit is mounted at the side of the runway and this can perform nearly the same function. Using intensity, rather than simple pressure, the direction of the noise can be tracked and thus aircraft waiting to take off at the 'HOLD" point can be removed from the data before acquisition. Because the Sound Gate in either pressure or intensity modes is not making an actual measurement, but simply providing a time and a rank order, there is no need to comply with acoustic standards. Thus, the residual pressure intensity index can be ignored as can the precise calibration and the response can be made to fit the exact parameters of the situation. The major parameters of the Sound Gate are exact time control and frequency stability. Thus, the levels referred to above are not dB(A) but frequency weighted in accordance with the expected noise.

This extra timing data contributes significantly to the system accuracy, as the unknown errors of timing are totally removed and the exact time of take-off or landing is known. This timing is now used instead of the data from the airport to correlate with the external noise monitors to give a good event definition and only then is the flight information added, giving good identification. Clearly, there have to be checks to ensure that the noise measured at an external monitor, the Sound Gate and the flight information are all relating to the same aircraft and new algorithms and hardware contribute to this.

5. SOUND GATE LOGIC

On a single runway airport, there are basically four possible situations with a Sound Gate system, as there can be a landing or a take-off in either of two directions. For simplicity, only one direction will be discussed, but the same logic is allied to the other direction. Consider a runway with Sound Gates (A) and (B) as shown in Figure 4.

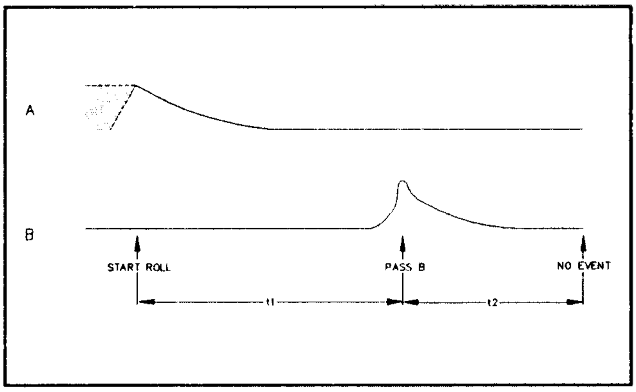

Figure 4: Take-off noise

In the upper trace for Sound Gate (A), there will be some unknown noise finally ending when the plane starts to roll. Until this time, the pilot may sit with his engine on high power, or he may simply coast onto the runway opening up the throttle when he is straight. Thus, in any event, there will be a high level which will die away as the plane goes down the runway. As the aircraft starts its run, the level at Sound Gate (B) starts to increase, reaching a maximum as the plane passes the Sound Gate itself. As the flight is in the air, the speed of passing is well controlled and the noise profile is reasonably defined no matter what the type.

The two Sound Gates send their data to a host computer which starts by taking the recognised data from (B) and searching the data file of (A) at the time it would expect a take-off and at the time it would expect the noise of a landing, i.e before and after its own event maximum. If it finds a falling slope inside pre-determined limits, at the relevant time from (A), it flags the event as a take-off. A similar situation occurs for a landing. Again, only the point where the aircraft is around touch-down is known. After touching down, the pilot may put on reverse thrust, simply throttle back, or even turn off the runway before reaching (A), so the output from (A) cannot be identified with any certainty in all cases. Once again, the start point of the correlation is the recognised shape at (B). In any particular airport, the traffic rules may insist that the aircraft continues to the end of the runway, or it may have to turn round and back-track. In these cases, the basic algorithm is modified to take account of these differences in operations. However, as the correlation is done on the host computer with clean data from each Sound Gate, the program is written in a high level language and thus can be easily modified for an individual situation without significant cost. This is done by having a standard set of libraries and taking the relevant data as needed.

The 'Sound Gate' host is a simple MS-DOS machine and it usually operates without even a monitor, as it simply works out the timing and passes this to the main host computer. This frees up the main host for its main task of reading the data from each noise monitor and storing the data in tables for subsequent use. Because of this low cpu usage, there is no need for complex operating systems such as UNIX and the whole program runs on a simple 486 machine with the windows (tm) operating system. Such a machine is about 20% of the cost of a UNIX device and the software is also far less complex and much more reliable in operation. Today, almost everyone can use MS-Windows, so the system requires very little in the way of operator training, a further bonus for the user.

6. SUMMARY

The new technique of adding a Sound Gate to an airport noise monitor can significantly reduce the errors in event timing and identifying of individual aircraft. The accuracy which can result, is likely to give rise to less than a single error per week at a busy airport, with correspondingly better hit rates at lower volume sites.

Because only sound is used to time and identify noise events, no external input to the system is required until the actual flight identification is to be correlated. This can then be done in delayed time once the real time noise data has been correlated as it happens.

Not only does this new method give better flight correlation than radar or passive listening systems alone, it is significantly cheaper and simpler to maintain over the equipment life.

Because the correlation is in a high level language, customised programs for individual airports are a simple task.

7. REFERENCES

1. A. D. Wallis & R. W. Krug "The Sydney and Brisbane Noise Terminals" Proc WESTPRAC pp 492-499 Nov 1991 Brisbane, Australia

2. United Kingdom Patent office Provisional patent.