Chapter 4

RESULTS

4.1 CHAPTER OBJECTIVES

The design and implementation of both the electronic and mechanical parts of the elephant recording collar were discussed in Chapter 3. In this chapter, the results of the evaluation of both the elephant collar and the elephant rumble detection algorithm are presented.

The general tests that have been done on the elephant collar are presented in Section 4.2 while the collar's field tests are discussed in Section 4.3. Section 4.3 discusses the results obtained by the automatic elephant rumble detection algorithm.

The microphone placement experiments were done to find the best possible way to attach the microphone to the collar. The stability of the electronics in extreme temperatures needed to be verified and the power consumption needed to be tested to ensure that it was within the specified range. Field tests were done to test the ability of the collar to operate under physically demanding circumstances.

Several aspects of the automatic elephant rumble detection algorithm were tested. The algorithm's ability to detect harmonic sounds in noisy conditions were quantified which lead to a better understanding of the algorithm's abilities.

4.2 GENERAL ELEPHANT COLLAR TESTS

The microphone placement experiments (4.2.1) and temperature stability and power consumption experiments (4.2.2) are described in this section.

4.2.1 Microphone placement experiments

The microphone of the elephant recording collar (i) needs to be sensitive, but (ii) needs to remain protected against physical abuse. The best solution appears to be to endeavour to mount the microphone in the electronics enclosure (which is strong and rigid), rather than distant from the enclosure, where it would not be protected as well. It was found that infrasonic sound does not propagate well through the 5 mm thick polycarbonate cover sheet. Several experiments were conducted to determine the best way in which to mount the microphone. These are described in this section.

The microphone was mounted in several different ways (described below) inside the enclosure and sound transmission properties were measured. To this end, several cover sheets were made and exchanged for each experiment.

4.2.1.1 Microphone mounting method 1

Seven holes with a diameter of 2.5 mm each were drilled in close proximity to each other (Figure 4.1). These holes do not go completely through the polycarbonate sheet, but end blindly 0.3 mm before the outer edge. This was done to maintain a smooth watertight surface on the outside of the cover.

4.2.1.2 Microphone mounting method 2

The microphone is mounted on the inside of the polycarbonate sheet directly beneath the holes. It is fixed to the cover panel with watertight Pratley putty. The remaining thick parts of the sheet between the holes provide strength in the area above the microphone.

Figure 4.1: Clear polycarbonate cover sheet shown from the outside, showing holes ending blindly.

Three holes with a diameter of 2.5 mm each were drilled through the polycarbonate sheet. A thin layer of softer plastic material (PVC, 1 mm thick) was glued (using Pratley clear glue) on top of the holes to provide a watertight smooth surface. The microphone was mounted directly beneath the holes.

4.2.1.3 Microphone mounting method 3

A single large hole with a diameter of 24 mm (so that the microphone can fit into it) was drilled into the polycarbonate sheet. Once again this hole does not go through, but ends blindly, leaving 0.3 mm of material on the outside of the sheet. The microphone is mounted inside this hole. See Figure 4.2.

Figure 4.2: This figure shows the microphone-diameter blind hole. The smaller holes at the bottom of the larger hole were drilled for method 5, but were not present in method 3.

4.2.1.4 Microphone mounting method 4

At first, sixteen holes with a diameter of 2.5 mm each were drilled through the polycarbonate sheet. These holes were filled with polyurethane foam to ensure watertightness. The microphone was placed directly beneath the foam filled holes. More holes were added in a later experiment (seen in Figure 4.3).

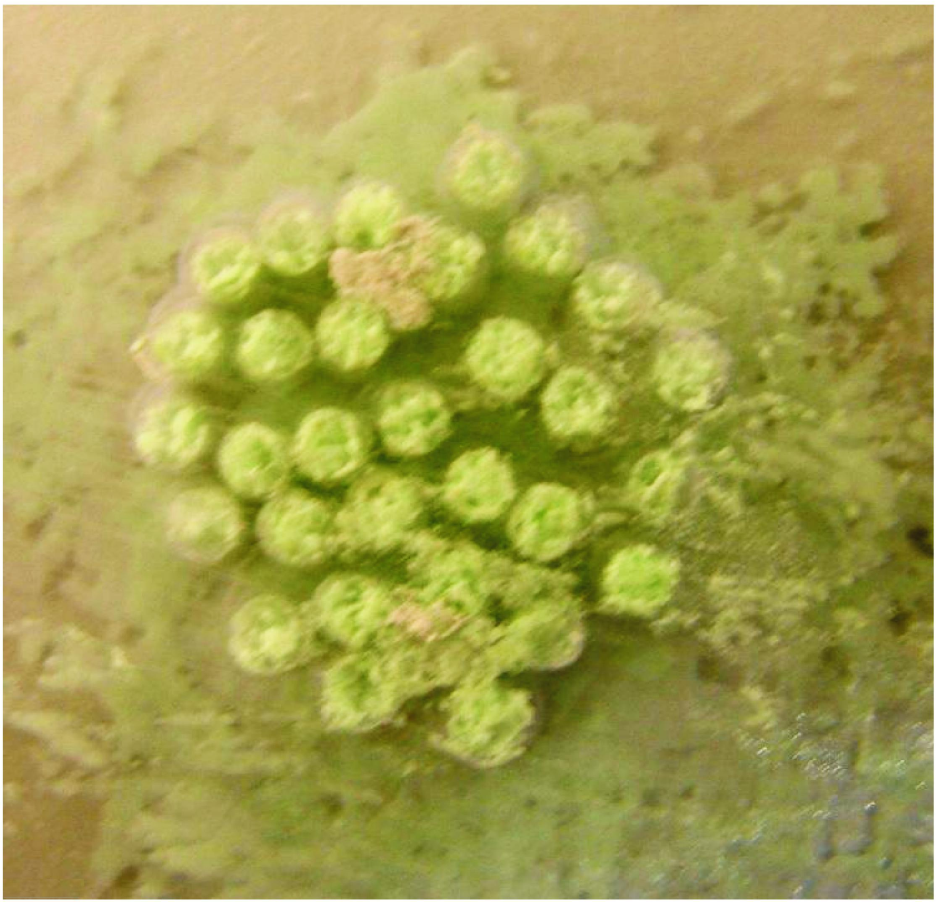

Figure 4.3: In this attempt, the holes were filled with polyurethane foam for a watertight seal.

4.2.1.5 Microphone mounting method 5

A microphone-sized hole as described in method 3 was drilled, again leaving a thin (0.3 mm) layer of material to cover the microphone. The microphone was wrapped in a thin sheet of latex to ensure watertightness and mounted into the hole. Twelve small holes with a diameter of approximately 0.4 mm each were drilled, aligned with the holes of the microphone casing (Figure 4.4).

Figure 4.4: The small holes drilled in the polycarbonate (left) are aligned with the microphone casing holes (right).

In this mounting method, the metallic enclosure of the microphone itself provides mechanical strength to the thin plastic area. The small diameter of the holes ensures that the surface of the polycarbonate sheet remains relatively smooth.

4.2.1.6 Procedure

Each of the mounting methods described above were experimentally tested as illustrated in the setup in Figure 4.5. A laptop running a custom Matlab program connected to a loudspeaker was used to generate a frequency sweep from 0 Hz up to 1250 Hz. The microphone was placed perpendicular to the loudspeaker at a distance of 500 mm. The output signal from the microphone was connected to a spectrum analyzer, which recording the transfer function.

4.2.1.7 Results and discussion

Figure 4.6 shows the results of the experiments. The first test was performed without a cover screwed onto the electronics enclosure, so that the transmission path to the

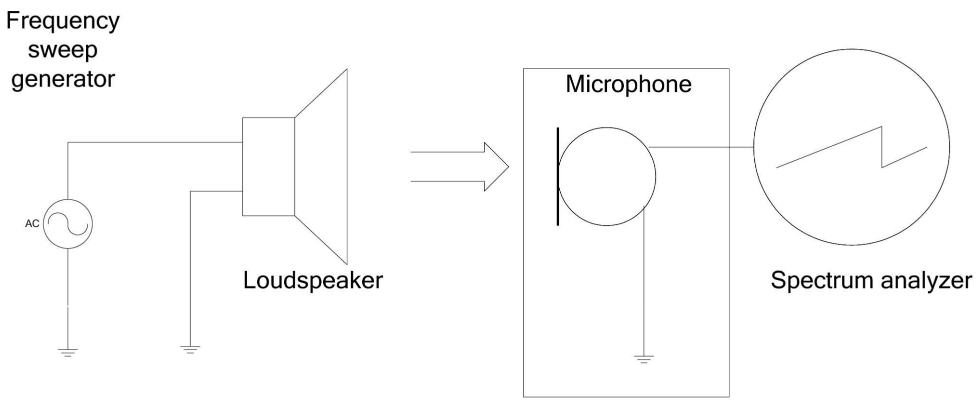

Figure 4.5: Experimental setup for measuring sound transfer functions.

microphone was unobstructed. This result was used as a reference for the various mounting methods tested.

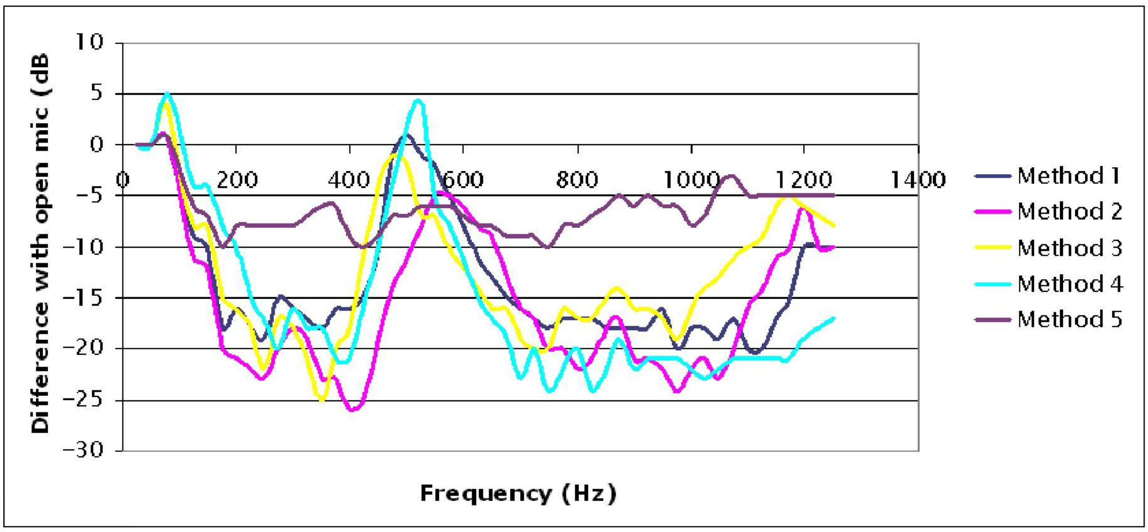

Figure 4.6: Transfer functions obtained from each of the five mounting methods described in the text, along with the reference ("open mic", i.e. unobstructed microphone).

The objective was to find a microphone mounting method that would provide good audio signal transmission, while providing adequate protection for the microphone. It is clear from Figure 4.6 that microphone placement methods 1 to 4 resulted in suppression of spectral components and appeared not to be appropriate choices. Mostly frequency components around 500 Hz conducted well through the thin layer of polycarbonate, but the frequencies lower and higher than 500 Hz were suppressed.

Figure 4.7 shows the difference between the reference microphone placement and each of the methods tested. It can be seen that methods 1-4 all have attenuation of more than 20 dB at some point in the desired band, while method 5 never attenuates more than 10 dB. Among all of the options, method 5 appears to be the most suitable.

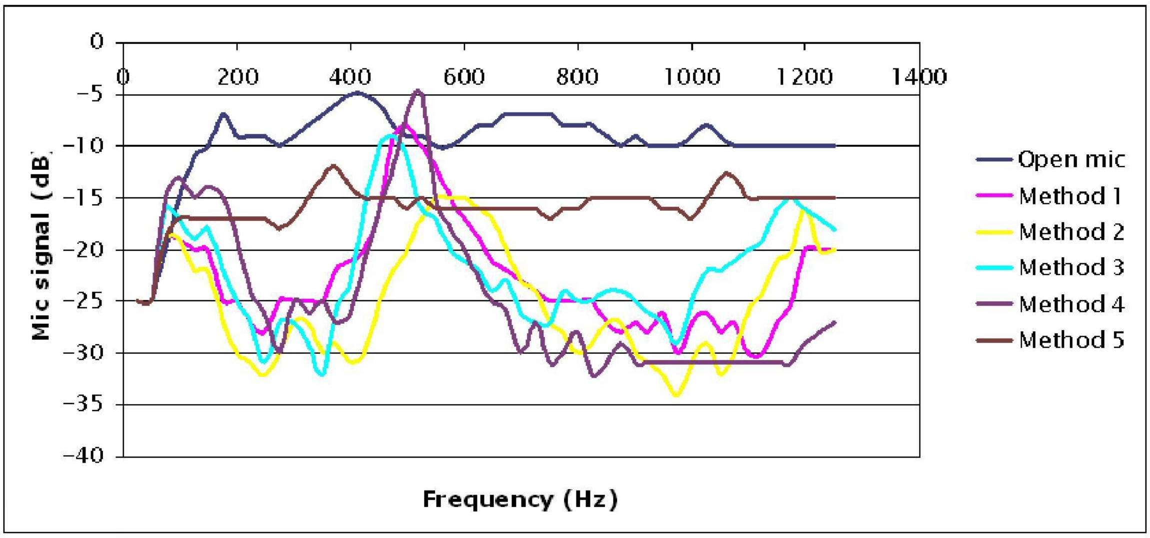

Figure 4.7: Difference (in dB) between the spectrograms of each of the above mentioned methods and that of the open microphone.

The results of the different methods were also compared in Table 4.1 using the average difference between an open microphone and all of the different options as well as the maximum difference.

Table 4.1: Comparison between microphone mounting methods in terms of attenuation

	Method 1	**Method 2**	**Method 3**	**Method 4**	**Method 5**
Average difference between open mic and tested method (dB)	-13.1	-15.32	-12.3	-14.98	-6.34
Max difference (dB)	21	27	29	27	11

Method 5 resulted in a flatter transfer function throughout the part of the frequency spectrum tested and appeared to be the better choice among these methods. Note also (although this result is not shown) that without the latex sheet covering the microphone in method 5, there was almost no attenuation. Thus, with holes in the cover panel carefully aligned with the microphone casing holes, the primary source of attenuation became the latex sheet used to provide a watertight seal.

4.2.2 Temperature stability and power consumption experiment

Simple experiments were conducted to determine the temperature stability of the device where the device (unpackaged and not mounted on the collar) was subjected to a range of different temperatures from -5 °C to 50 °C. These temperatures probably exceed the temperature extremes to which the recording collar would be subjected in the field. The current consumption of the device was measured at each temperature step to determine the influence of temperature variations on power dissipation. Recorded data was verified after each step to determine that correct operation of the device was maintained at every temperature level.

4.2.2.1 Objective of experiments

The objective was to verify that the elephant recording collar would operate correctly at the extreme maximum and minimum temperatures that it might encounter when operating in the field.

4.2.2.2 Equipment

The equipment used in the temperature stability tests is listed below:

1. The elephant recording collar electronics
2. The actual battery pack that would be used in the field
3. A digital multimeter
4. A freezer
5. An oven

6. A personal computer

7. A CompactFlash Card reader

4.2.2.3 Method

The experiment was set up as shown in Figure 4.8. The battery pack was connected to the electronics through an ammeter to measure the current drain of the system. A continuous pure tone at 300 Hz was played through a loudspeaker and recorded (via the microphone) by the elephant recording collar electronics. The device then wrote recorded data to the CF card. After completion of recording, the CF card could be read by a PC using a CF card reader.

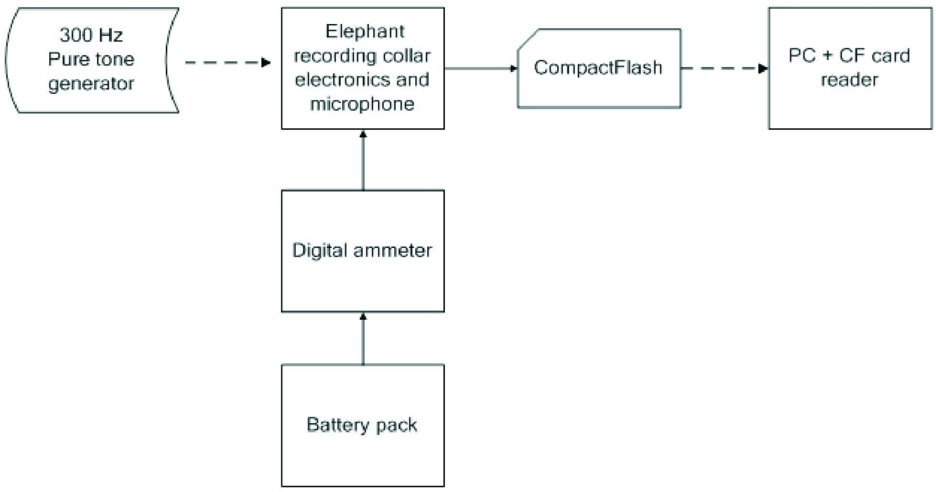

Figure 4.8: Setup of the experiment.

4.2.2.4 Procedure

1. The setup as shown in Figure 4.8 was put into the freezer (set at -10 °C) for 30 minutes.

2. At the end of the 30 minutes the current drain was measured.

3. The files on the CF card were transferred to the PC.

4. It was verified that correct operation had occurred throughout the experiment by checking that the files contained the correct sound data throughout the 30-minute recording. "Correct operation" is defined as no errors having occurred in the recorded 30-minute data file.

5. The same procedure was repeated in an oven that had been pre-heated to 50 °C.

6. These results were then compared to previous measurements at room temperature.

7. The experiment was repeated only once at each temperature.

4.2.2.5 Results

Table 4.2 shows the results obtained from the experiments.

Table 4.2: Operation at different temperatures

Temperature (°C)	**Current consumption**	**Correct operation**
-10	19 mA	Yes
25 (room temperature)	19 mA	Yes
50	21 mA	Yes

4.2.2.6 Discussion

The effects of temperature on the accuracy of the measuring apparatus appears to be negligible. The results of the experiments suggest that the electronics would be able to operate correctly between the two temperature ranges (-10 to 50 °C) as examined in this experiment. The current drain of the device was approximately 20 mA, which was within the range allowed by the specifications.

4.2.3 Field tests

The prototype elephant recording collar was tested on a young elephant bull for one week in August 2007. The Elephant Sanctuary near the Hartebeespoort dam allowed us to test the collar on one of their elephants. These elephants are tame enough to allow handlers to put the collar on and take it off without the need to tranquillize the elephant. This was ideal for the first brief field tests to sort out any obvious problems that might arise. Figure 4.9 shows the handlers commanding the elephant to lie down so that they could attach the collar.

Figure 4.9: The elephant lying down so that the handlers may mount the recording collar.

4.2.3.1 Method

A 4 GB micro hard drive with a CompactFlash interface (from now on called the microdrive) as well as a 256MB CompactFlash card were used as data storage media. The cards were inserted into two of the available CompactFlash sockets housed in the collar. The device was activated approximately 15 minutes before the collar was mounted on the elephant. Figure 4.10 shows the collar fitted onto the elephant.

Figure 4.10: The elephant after the collar have been successfully mounted.

4.2.3.2 Results

After one week the collar was retrieved. The electronics were still in working condition and it appeared that no dust or water entered the sealed part of the collar. This was probably not an adequate test for watertightness, since elephants apparently do not take mud baths in winter. The polycarbonate cover of the collar was still firmly attached and appeared to be an appropriate choice of material for this application.

The data was retrieved from the cards by putting the card into a card reader that was connected to a windows operated computer. The computer detected the card as a removable drive with a FAT32 file system as is shown in Figure 4.11. Sound files are in the wave file format. Each sound file has a duration of one hour and starts where the previous file ended, resulting in the availability of continuous sound data. Each sound file has a corresponding text file that contains the time and date, temperature and GPS coordinates of the collar as it was at the time of file creation.

Inspection of the data on the memory card revealed that the device stopped recording 30 hours after activation and returned to the reset mode. The 4GB microdrive was

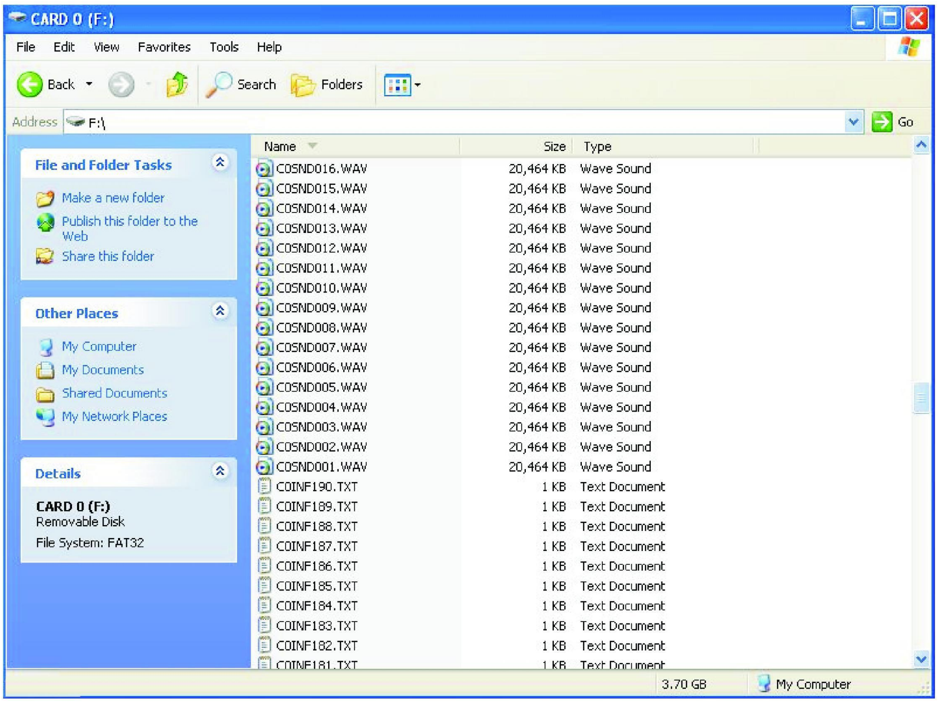

Figure 4.11: The root folder of a CF card containing the recorded sound and information files.

the only card to contain data as the recording process stopped before the 256 MB card could be reached. At first it was thought that power from the battery could have been lost for an instant (as a result of a bump) causing the device to restart and return to reset mode. Further investigation revealed that the specific microdrive that was used consumed more energy than expected. The fact that all previous testing of the collar was done with the same battery pack and microdrive resulted in the battery pack running flat before the calculation recording time had expired. This meant that the device was still entirely operational and only needed new batteries.

This first field test highlighted some problems, but also, despite the disappointingly short recording time, showed that good quality recordings of elephant calls could be made. Some examples are shown here. The rumbles shown in the spectrogram of Figure 4.12 were recorded on file C0SND011.WAV during the field test. Figure 4.13 also shows a spectrogram containing elephant rumbles from file C0SND024.WAV.

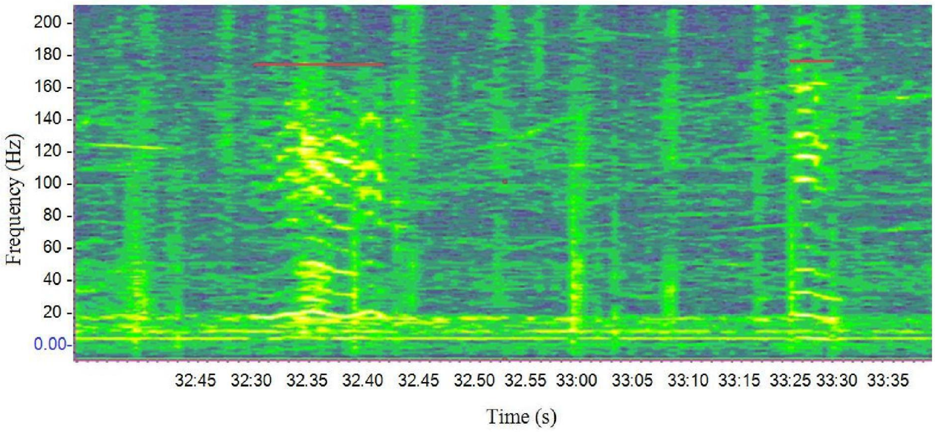

Figure 4.12: A spectrogram from a part of the file C0SND011.WAV recorded during a field test containing elephant rumbles.

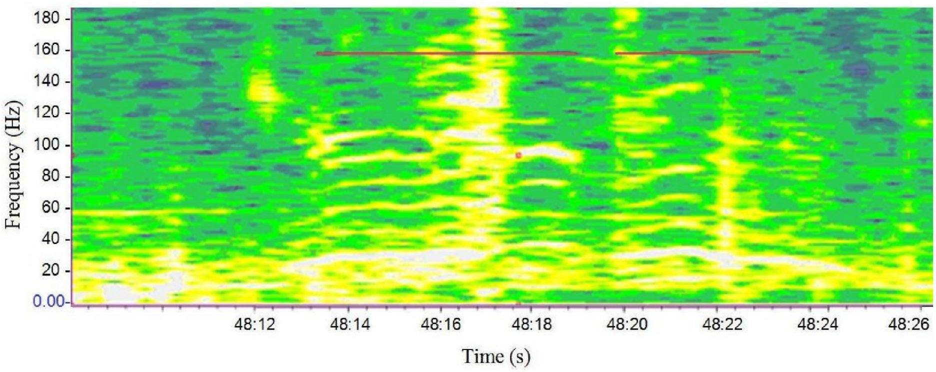

Figure 4.13: A spectrogram from a part of the file C0SND024.WAV recorded during a field test containing elephant rumbles.

4.2.3.3 Apparent errors in the recorded data files

It appeared that sound was not properly recorded in the first few files (where the battery pack was still operational), while recordings in later files were of good quality. So far, we have not found a good explanation for this. Because the microphone was mounted in an airtight container, one possible explanation may be the difference in air pressure between the location where the collar was stored and where it was activated, resulting in the membrane of the microphone being "sucked in" for a while until the pressure could equalize.

A periodic clicking sound occurred in some of the recorded sound files about every 20 seconds as is the case shown in the spectrogram of Figure 4.14. This was not noted in previous tests in the lab. The time interval of the clicking sounds corresponded exactly with the buffer size used in this prototype device. There did not appear to be any previously unnoticed buffer overflow and when using a new power supply the clicking sounds ceased. It appeared that, because the energy levels of the battery were low, the high spike current drawn by the microdrive during the writing cycle (occurring each time that a buffer needed to be written to permanent storage) caused the supply voltage provided by the battery to dip, in turn causing the voltage reference point of the ADC to fluctuate briefly and thus causing the single spike in the sound. Thus, this error was also caused by battery failure.

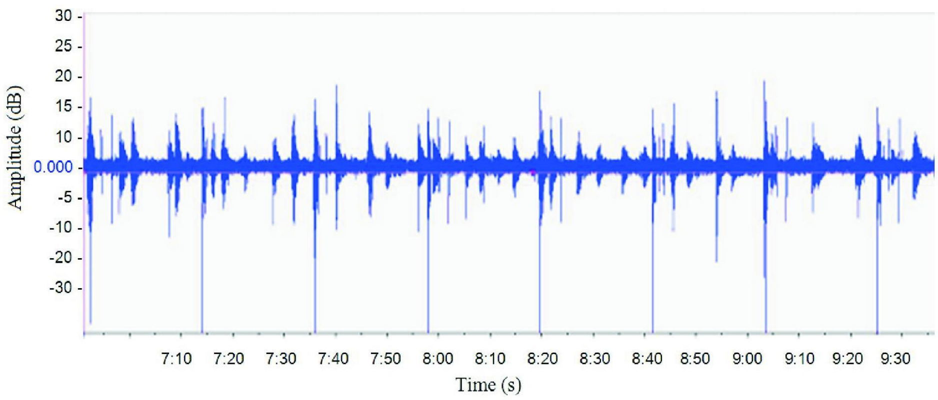

Figure 4.14: The recorded sound contained periodic clicking sounds.

One of the sound files contained a periodic sound that initially appeared to be another error (Figure 4.15).

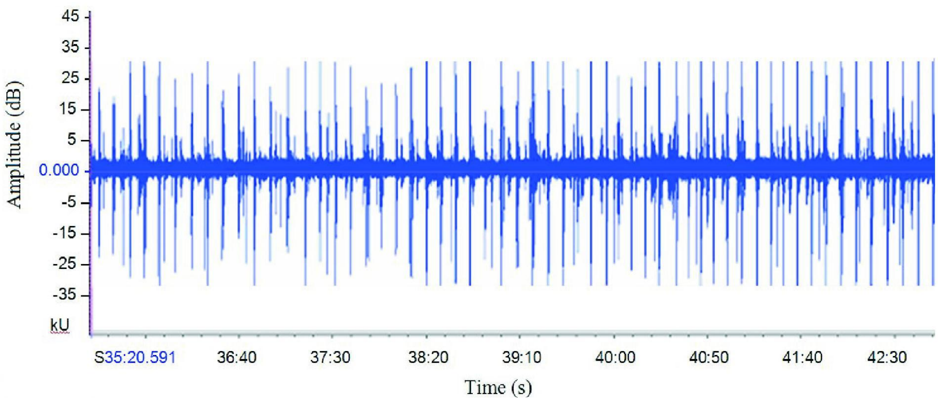

Figure 4.15: (Almost) periodically occurring sounds were recorded.

However, closer investigation suggested that the periodic sounds occurring in wavefile C0SND011.WAV as shown in Figure 4.15 was not a design defect. We could not repeat this effect in the lab, and the almost (but not entirely) periodic sound suggested that this was in fact a sound captured by the microphone. When the recording was played 10 times faster than normal, a sound resembling a person taking quick successive breaths could be heard. Given the time of this recording (middle of the night), it might be possible that a snoring elephant was recorded.

4.2.3.4 Second field test

A second field test was done with a fresh battery pack. The test was done at the same location and with the same elephant that was used for the first test. This time around a 4 GB CompactFlash card was used instead of the microdrive. The test once again had a duration of one week.

The following results were obtained:

1. The test was completed successfully. Both the memory cards were filled to capacity with recorded data.
2. The electronics were still in working condition. Plenty of rain had fallen during the test, but it appeared that no dust or water entered the sealed part of the collar.

3. It was noticed that, as expected, the defect that was observed in the first field test was not present in these recordings.

4. Some of the recordings seemed to be recorded at a softer sound level than the rest of the recordings.

4.2.3.5 Conclusion

The results of the first field test confirmed the ability of the elephant collar to record elephant vocalizations in realistic circumstances. The mechanical design of the collar proved able to withstand the physical conditions that it had been subjected to while fitted on the elephant.

The recording process ended prematurely because of earlier than expected end of battery life. This was the result of the microdrive having had a higher power consumption than expected. A defect was observed in some of the recordings. A phase shift occurred in the sound every time that the information in the RAM buffer was transferred to the microdrive. This happened because the excessive power consumption of the microdrive disturbed the reference voltage of the ADC at low battery capacity.

A second field test was done with a fresh battery pack and a Compactflash card was used instead of the microdrive. This solved the problems that had been encountered in the first field test, since the recording process was completed successfully and without the presence of the clicking sound in the recordings.

The fact that some of the recordings were at a softer sound level was probably due to mud blocking the microphone for the period of time that it needed to dry and peel off.

4.3 SIGNAL PROCESSING RESULTS

Finding a measure of the success of the elephant rumble detection algorithm is a difficult task because it can only be referenced against the subjective findings of a person who did a visual analysis of the spectrograms of the recordings. An experiment that aims to quantify the abilities of the algorithm in an objective way will be presented. This experiment is discussed in Section 4.3.1.

The next step was to test the algorithm on actual elephant vocalization data to see if the results correspond to the findings of the visual identification of the same data (4.3.2). A number of good quality elephant rumbles were isolated from the available elephant recordings by visual inspection of spectrograms. The algorithm was tested on clearly defined rumbles as well as rumbles that occurred together with common background noises. The results were used to do an estimation of the accuracy with which the algorithm could detect elephant rumbles (4.3.3).

Finally, the circumstances under which the algorithm will fail (4.3.4) will be investigated and a method by which the algorithm can detect overlapping rumbles will be discussed (4.3.5).

4.3.1 Quantifying the pitch detection and tracking ability of harmonic sounds in noisy conditions

It was not an easy task to measure the level of success of the elephant rumble detection algorithm since the results of the algorithm had to be compared to the subjective findings of a person who did a visual analysis of the spectrograms of the elephant recordings. It was therefore decided that an experiment should first be done to quantify the performance of the algorithm based on certain parameters that could be measured and controlled.

The algorithm used the fact that elephant rumbles have a harmonic structure to distinguish between rumbles and other noises present in the recording. The following experiment generated sounds with a varying number of upper harmonics. White noise (like light wind induced noise often found in elephant recordings) was added to the signal. The energy of the white noise was increased repeatedly until the algorithm started to fail (tracking of the pitch was lost). This were repeated with other sounds with upper harmonics.

4.3.1.1 Method

This experiment was carried out using the Matlab programming language. A pure tone signal with amplitude 1 and frequency 25 Hz was generated and white Gaussian noise with zero mean was added to the signal. In this case "white" implies noise with a flat PSD (Power Spectral Density) from 0 Hz to 1500 Hz. The power of the noise signal was increased until some gaps in the pitch track started to occur. The SNR at which this happened was recorded.

The same procedure was then repeated, but with a first overtone added to the signal. The signal now consisted of the 25 Hz signal and a 50 Hz signal also with an amplitude of 1. Once again the noise power was increased until some gaps in the pitch track started to occur. The SNR at which this happened was recorded.

The same procedure was repeated, each time with one more harmonic component added to the signal.

4.3.1.2 Results

Table 4.3 shows the results obtained from the experiments.

Table 4.3: Results obtained by the SNR experiment

SNR (dB)	25	6	-5	-8	-8.4	-8.7
Number of overtones	0	1	2	3	4	5

Figure 4.16 shows a plot of the SNR values obtained in the procedure explained above. The figure shows that the pitch of a signal containing more than one overtone can be tracked in noisy conditions.

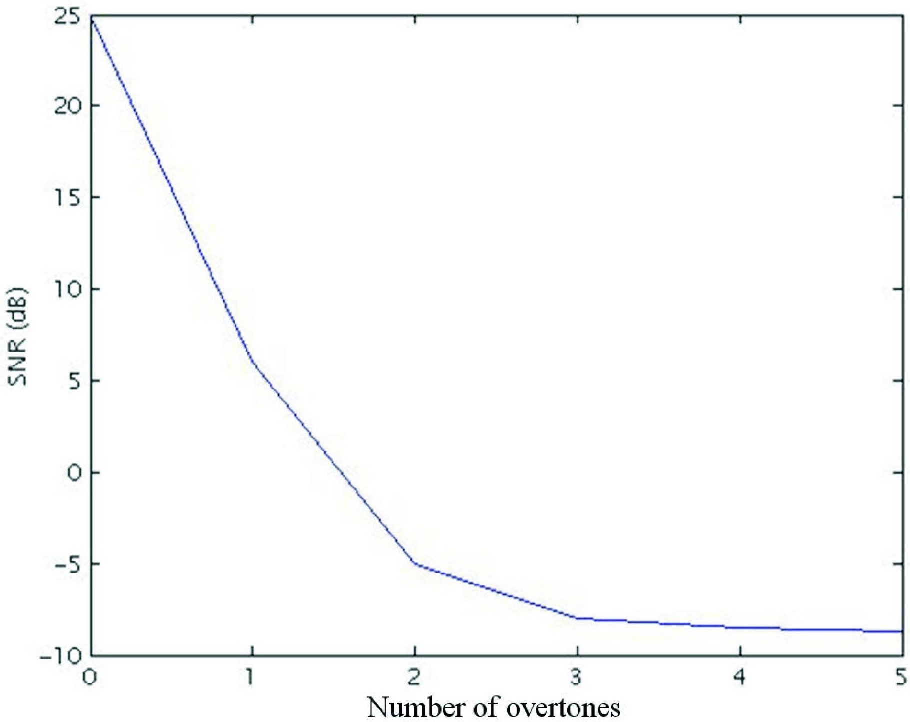

Figure 4.16: A plot showing the minimum SNR values at which a continuous pitch track could be estimated versus the number of harmonic frequencies present in the signal.

4.3.1.3 Discussion

The results indicate that the pitch of a signal which consists of only the fundamental frequency (has no harmonics) will not be tracked if some noise is present in the recording. This is a useful quality, as it can assist in avoiding false alarms generated by narrow band noise that could be produced by motor vehicle engines (as an example see the first 30 seconds of the spectrogram shown in Figure 4.17).

The results also show that the algorithm will work best when the elephant rumbles in a recording have three or more clear overtones. The algorithm can detect the pitch of a sound with harmonic structures in an SNR as low as -8 dB which means that the algorithm has good rejection of broadband noise like that caused by wind. This result also implicates that the algorithm's weakness lies in its inability to reject noises with a harmonic structure.

4.3.2 Tests on recording segments containing single elephant rumbles

As stated in Chapter 3, a number of elephant recordings made in the Kruger National Park were made available for this study. As a first step in testing the rumble detection algorithm on actual elephant recordings, some clearly definable elephant vocalizations were located within the provided recordings. The segments of the recording containing the vocalization were saved as a separate sound file. A spectrogram of one such sound segment (R001.WAV) is shown in Figure 4.17. The part of the spectrogram where the rumble occurred was marked by a red line. It can be seen that the rumble shown on the spectrogram started at approximately two and a half seconds into the recording and ended at approximately 5 seconds.

The sound segments were then used as input to the elephant rumble detection algorithm. It was found that algorithm accurately detected an elephant rumble in a good quality sound segment (if the rumble is clearly definable and little background noise is present). The output produced by the algorithm is shown in Figure 4.18.

Certain rumble parameters were automatically determined by the algorithm and are given in Table 4.4. The time in the sound segment where the presence of an elephant

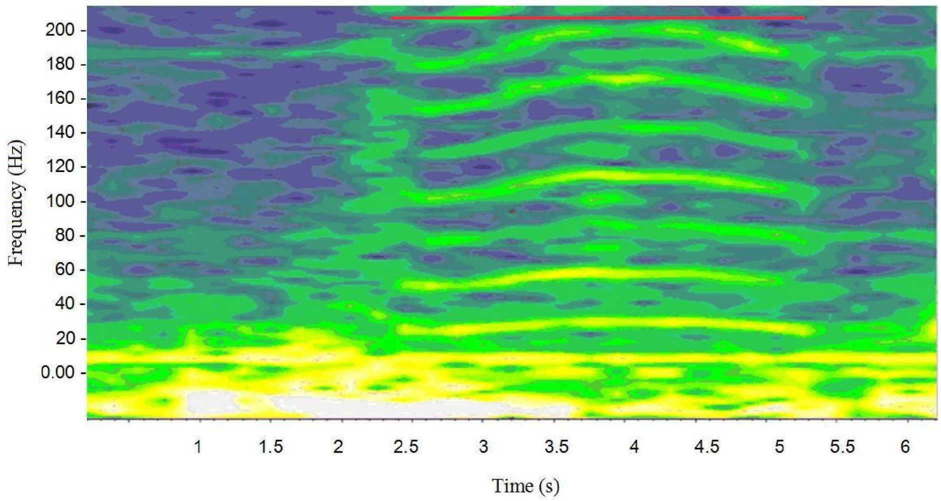

Figure 4.17: Spectrogram of a sound segment containing a clearly definable infrasonic elephant rumble.

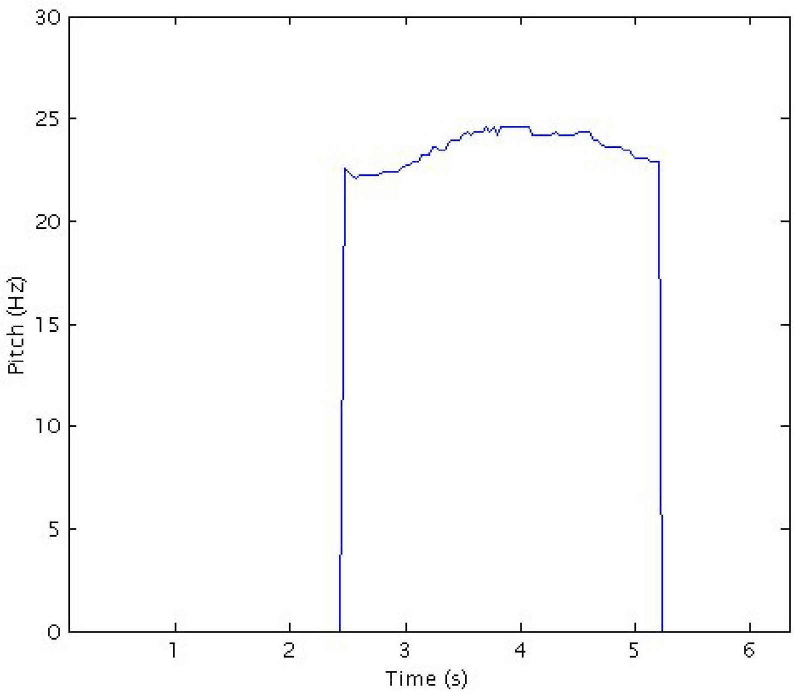

Figure 4.18: Output from the rumble detection algorithm after processing the sound segment of Figure 4.18.

rumble was first detected is denoted as *time begin* and the value is given in seconds. *Time end* indicates the time in the recording where the vocalization stopped. The duration of the call was calculated by subtracting the end time from the starting time of the call. *Pitch high* denotes the highest pitch present in the rumble and is given in Hz. *Pitch low* indicates the lowest pitch that was detected in the rumble. *Pitch average* gives the average pitch of the rumble and is calculated by adding the pitch value of all the time windows within the rumble together and dividing the answer by the number of time windows present in the rumble.

The results obtained by the automatic detection algorithm correspond well with the visual detection of the rumble from the sound segments. The automatic calculation of the pitch values is something that cannot be done by visual inspection of a spectrogram and can be a valuable asset since the pitch of a rumble can communicate information about the emotional state of an elephant (Soltis *et al.*, 2005b).

Table 4.4: Automatically detected parameters of the elephant rumble shown in Figure 4.17

Time (s)		**Call Duration (s)**	**Pitch (Hz)**		
Begin	**End**		**High**	**Low**	**Average**
2.4	5.2	2.8	24.8	22.3	23.6

4.3.3 Tests on recordings with background noise

It was established that the rumble detection algorithm could accurately detect the location of a clearly definable rumble within a sound segment. For the algorithm to be used in a meaningful way, it would have to be able to determine rumble locations within recordings containing common background noises. A few such examples will now be given.

These examples were chosen from approximately 40 hours of available data. Notes (compiled by elephant call experts) were provided with the data. These documented the conditions under which the data were recorded, as well as the distance from the elephant and a rough estimate of when vocalizations were uttered during the recording. After careful analysis of some of the data, it became clear that, in general, no calls could be detected in recordings made from too far away from the elephants. Further, on extremely windy days the data were corrupted. All recordings were made on two channels.The recordings on the second channel could not be used because the sound level was too low. Where recordings were made under the unwanted conditions mentioned above, these were disregarded, resulting in roughly a tenth of the data being analysed.

The data used in the following examples were recorded in typical conditions, defined here to be from a location close to the elephants, relatively far away from roads and not in strong wind. From preliminary work, it was known that the algorithm should perform well if the harmonic structure of the rumbles were well preserved and if the recording was not contaminated with unwanted noises that had strong harmonic structure. Circumstances under which the algorithm failed are discussed in Section 4.3.4.2.

The spectrogram of the sound (S0017.WAV) used to demonstrate the manual identification of elephant rumbles in Chapter 3 is shown in Figure 4.19. The elephant rumbles that were manually identified are indicated with red lines. The horizontal green lines that can be seen in the spectrogram are unwanted narrow band noise produced by a car engine. The vertical green lines are broadband noise typically caused by wind, breaking of brances or disturbance of the microphone itself.

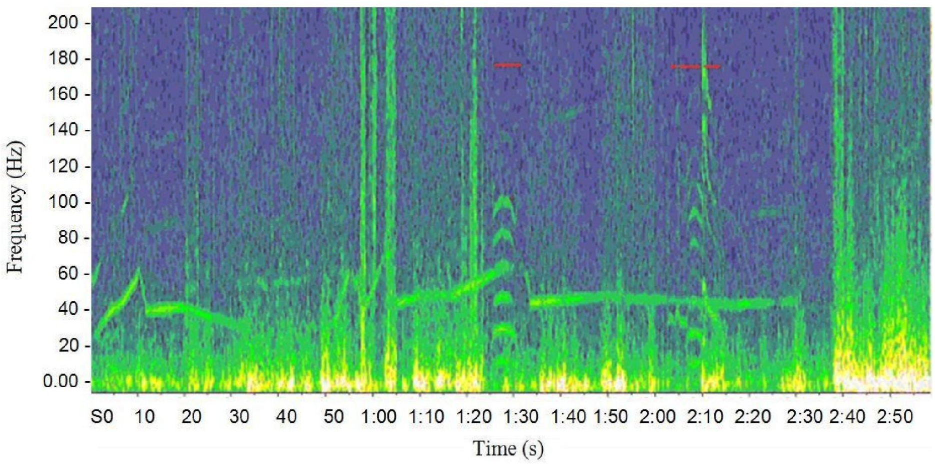

Figure 4.19: A spectrogram of a segment of sound containing narrow band and broadband noise as well as elephant rumbles (which are indicated by red lines).

Figure 4.20 shows the results obtained after processing the recording with the algorithm. It can be seen that unwanted noises are rejected well, while the harmonic structures of the elephant calls are all identified and their pitch extracted.

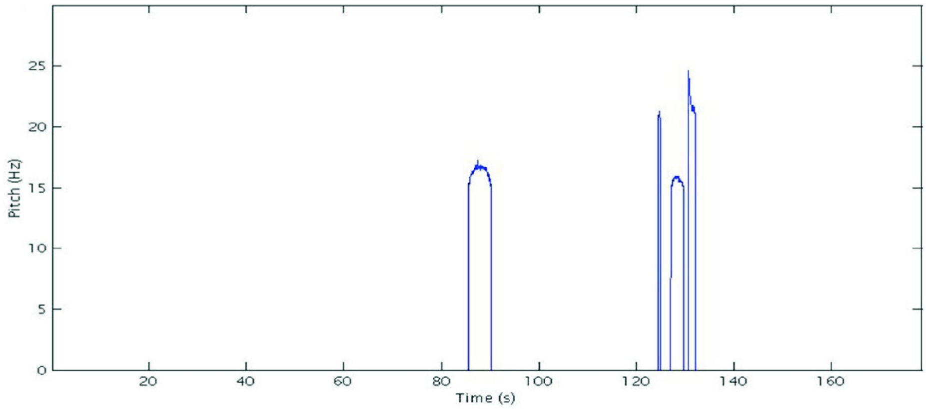

Figure 4.20: Final output of the algorithm after the file in Figure 4.19 was analysed.

The spectrogram shown in Figure 4.21 shows a more difficult example (S0023.WAV). Engine noises can once again be seen in the middle of the recording. Note that certain parts of the engine noises contain upper harmonics. It seems that this recording was made from a considerable distance from the elephants, since the rumbles that were recorded are quite weak. The rumbles that were manually identified are indicated by a red line.

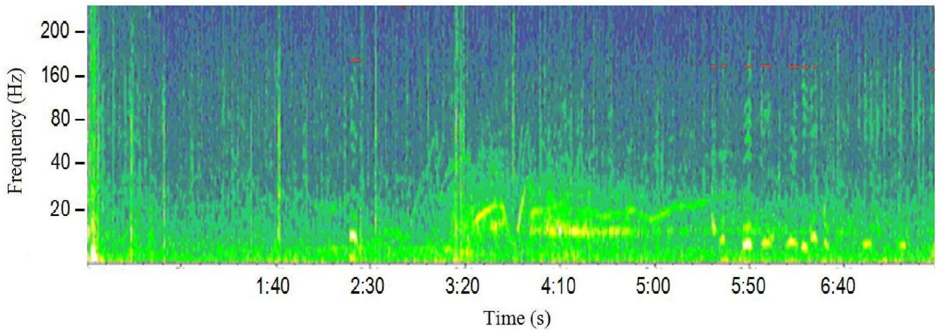

Figure 4.21: A spectrogram of a segment of sound containing engine noise as well as elephant rumbles.

Figure 4.22 shows the results obtained by the algorithm after the processing of the sound in Figure 4.21. The noise was successfully rejected. The harmonic engine sounds were rejected because the rate of change of the pitch was high enough that a pitch track could not be formed. The fifth manually detected rumble was not detected by the algorithm due to the loss of a clear harmonic structure.

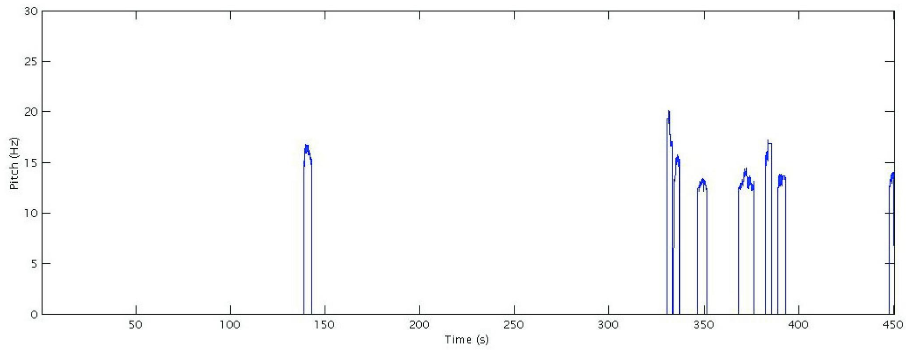

Figure 4.22: Final output of the algorithm after the file in Figure 4.21 was analysed.

Figure 4.23 shows the spectrogram of a recording (S0039.WAV) containing engine noises and very loud screaming type elephant sounds. When listening to the recording speeded up ten times faster than normal, it could be heard that rumbling sounds were also made together with the screaming sounds. These rumbles were very faint and were not easy to identify visually. The rumbling sounds that were manually identified are indicated by red lines.

Figure 4.24 gives the results produced by the algorithm. Noises were rejected successfully, but one of the calls that was identified manually was not detected due to a lack of a sufficient number of higher harmonics.

Table 4.5 shows a list of the calls (presented in this section) that have been identified manually along with the results of the automatic rumble detection algorithm. It indicates whether the algorithm was able to identify the rumble and presents the parameters obtained by the algorithm. A rumble was labelled as correctly identified when the start and end times were within 0.3 seconds of the call identified manually from the spectrogram. The manually identified rumble locations have indistinct starting points and end points and are for this reason not given in the table.

The results that were obtained by the elephant rumble detection algorithm are given in Table 4.5. The recordings that were analysed up to this point were chosen to be representative of the average quality of useful elephant recordings as was discussed in

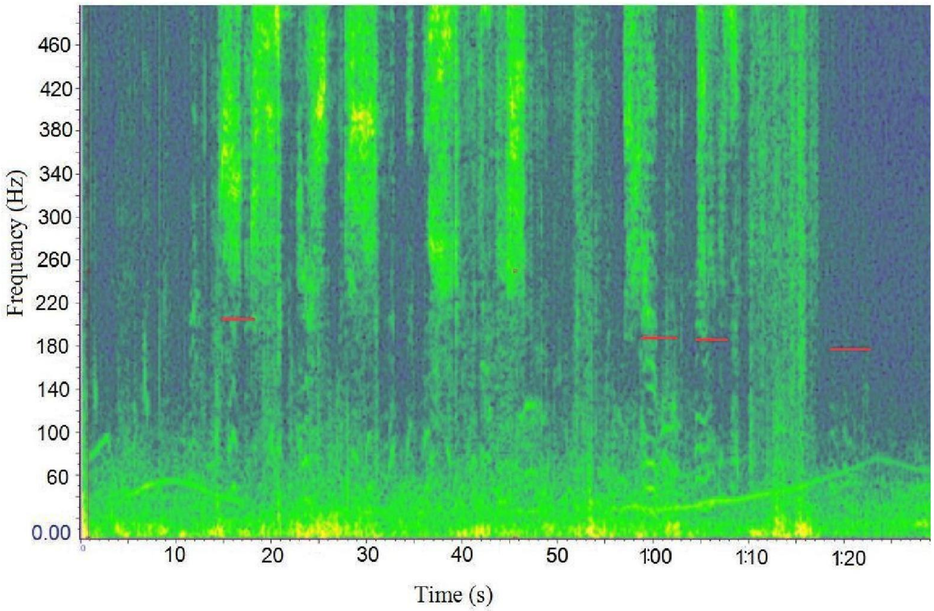

Figure 4.23: A spectrogram of a segment of sound containing scream-like elephant sounds as well as elephant rumbles.

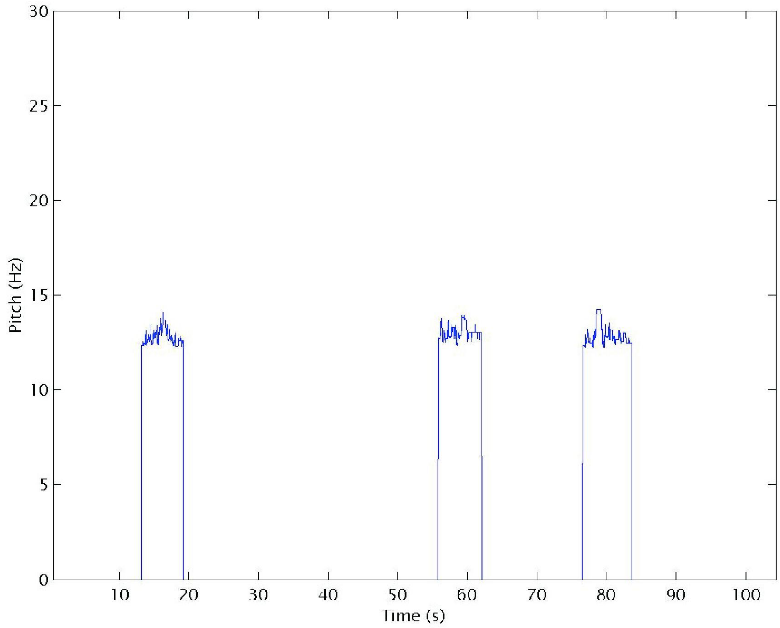

Figure 4.24: Final output of the algorithm after the recording of Figure 4.23 was analysed.

the opening paragraphs of this section. All but two of the calls that were identified by hand were also correctly identified by the algorithm. Four clearly defined elephant rumbles including the one depicted in Figure 4.17 were isolated and were correctly identified by the algorithm. Combining these with the results given in Table 4.5 it can be seen that 19 out of the 21 calls that were manually detected were also detected by the algorithm. This means that the number of calls that were correctly identified by the elephant rumble detection algorithm corresponded with 90.47% of those that were identified manually.

4.3.4 Circumstances under which the algorithm fails

Within some limitations, the algorithm discussed here proves to be a reliable way of detecting infrasonic elephant rumbles in noisy conditions. It also rejected unwanted low frequency sounds that do not have higher harmonics, such as distant motor vehicles. Sounds that contain higher harmonics, but had a rapid rate of change in pitch were also rejected.

Except for the examples shown in the previous subsection, where faint calls were not detected because of the masking of the harmonic structure by noise, there are certain other instances where the algorithm might fail.

4.3.4.1 Overlapping calls

Figure 4.25 shows the spectrogram of a sound (R005.WAV) where a faint rumble appears in the background, while a short louder rumble occurred simultaneously.

Figure 4.26 shows that only the louder call was identified by the algorithm. It will be shown in the next subsection that overlapping calls can be detected by applying a slight modification to the algorithm.

Table 4.5: A summary of the results obtained from the elephant rumble detection algorithm.

Recording: Figure 4.19		**Data produced by the algorithm (S0017.WAV)**					
Call No.	**Algorithm identified?**	**Time (s)**		**Call Duration (s)**	**Pitch (Hz)**		
		Begin	**End**		**High**	**Low**	**Average**
1	Yes	85.6	90.1	4.5	17.6	15.1	16.8
2	Yes	124.4	125	0.6	22.4	21.9	22.1
3	Yes	127.1	129.8	2.7	16.1	14.7	15.3
4	Yes	130.7	132.1	1.4	24.6	21.8	22.3
Recording: Figure 4.21		**Data produced by the algorithm (S0023.WAV)**					
Call No.	**Algorithm identified?**	**Time (s)**		**Call Duration (s)**	**Pitch (Hz)**		
		Begin	**End**		**High**	**Low**	**Average**
1	Yes	139	143.1	4.1	17.0	14.9	15.4
2	Yes	330.8	333.4	2.6	20.0	16.8	17.1
3	Yes	334.4	337.1	2.7	14.9	12.7	13.8
4	Yes	346.3	351.4	5.1	13.3	12.5	12.9
5	No	N/A	N/A	N/A	N/A	N/A	N/A
6	Yes	368.4	376.5	8.1	14.8	12.6	13.0
7	Yes	382.5	385.7	3.2	17.4	14.3	16.2
8	Yes	389	392.9	3.9	14.2	12.3	13.7
9	Yes	447.6	450.2	2.6	14.1	12.7	13.5
Recording: Figure 4.23		**Data produced by the algorithm (S0039.WAV)**					
Call No.	**Algorithm identified?**	**Time (s)**		**Call Duration (s)**	**Pitch (Hz)**		
		Begin	**End**		**High**	**Low**	**Average**
1	Yes	12.6	19.7	7.1	14.3	12.4	12.9
2	Yes	56.8	62.3	5.5	14.1	12.6	13.1
3	No	N/A	N/A	N/A	N/A	N/A	N/A
4	Yes	77.2	83.5	6.3	14.9	12.3	13.2

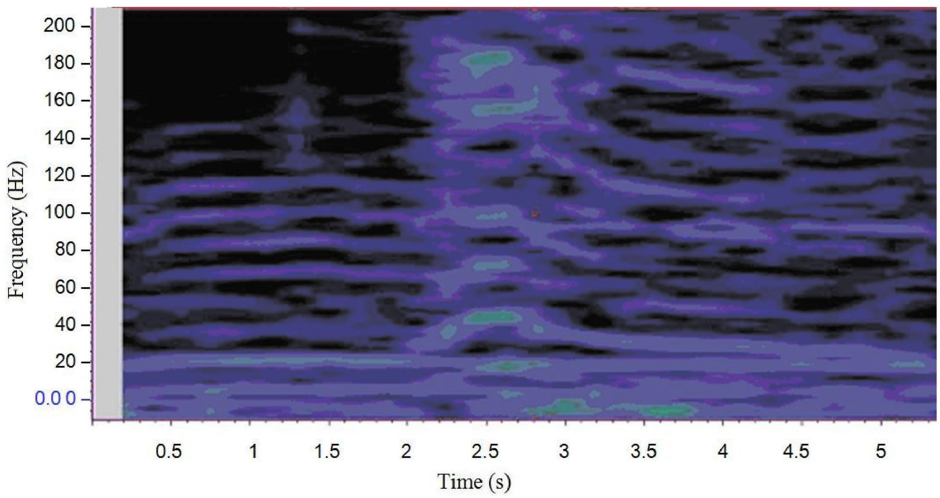

Figure 4.25: Spectrogram showing a faint call in the background with a louder call starting in the middle of the spectrogram.

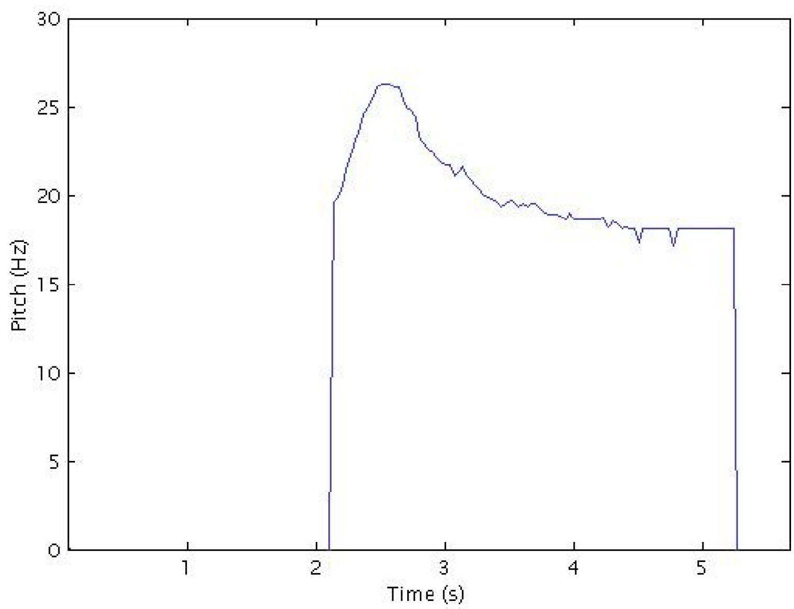

Figure 4.26: The output of the algorithm after the sound in Figure 4.25 was processed.

4.3.4.2 Unwanted harmonic noises

Figure 4.27 shows the spectrogram of a recording (003.1001.6000.WAV) with very energetic unwanted frequency components at 50 Hz and 75 Hz and the engine noises of a motor vehicle. Two elephant rumbles occur at 3:20 and 3:30, but it can be seen from the spectrogram that the energy of the unwanted components are much stronger than those of the rumbles itself. The algorithm falsely detected the engine noises as rumbles (because of the presence of the stronger harmonics) and at some point detected the pitch of the unwanted frequency components for a short while. This can be seen in Figure 4.28 which shows the output of the algorithm when the sound shown in Figure 4.27 was the input. The last of the two elephant rumbles was less faint and was correctly identified.

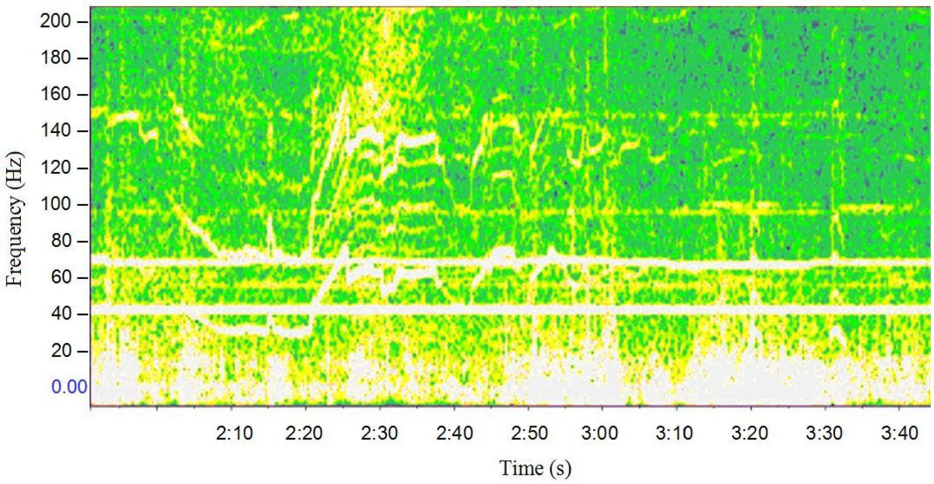

Figure 4.27: Spectrogram of a recording with loud unwanted frequency components.

The quality of the recording of Figure 4.27 is very low. The sources of the unwanted frequency components seems to be very close to the microphone while the elephant calls were distant. Although a recording as the one shown in this example is unlikely if an elephant recording collar was used, it is a good example to use for pointing out the weaknesses of the elephant rumble detection algorithm.

An unwanted sound with strong harmonics in the infrasonic bandwidth causes false rumble detections, while distant elephant rumbles which lost their upper harmonics might not be detected.

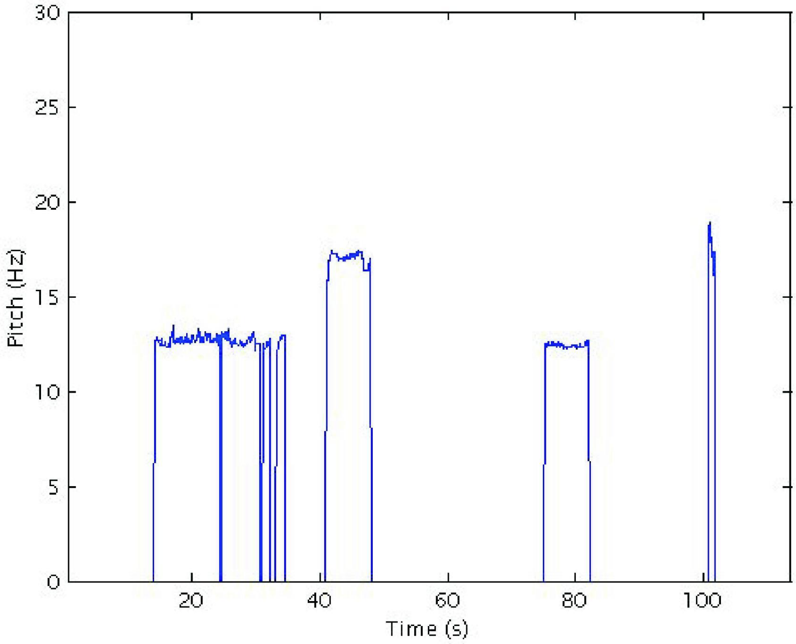

Figure 4.28: The algorithm's output after the sound shown in Figure 4.27 was processed.

4.3.5 Detection of overlapping rumbles

In an attempt to mitigate the problem of detecting overlapping calls, the algorithm was adapted slightly to allow tracking of two pitch tracks simultaneously. In the single-track algorithm, the point in the summed correlogram where the maximum value occurred was used to calculate each data point of the pitch information array. In the adapted version, the two peaks in each summed correlation with the highest values were used to generate two pitch information arrays. The information contained in these two arrays was then processed in the same way as in the case of the single pitch information array with the exception that two pitch tracks were processed. In essence, the procedure described in Section 3.5.2.6 is performed twice.

Figure 4.29 shows an example of the graphical output generated by the dual pitch track algorithm. The sound shown in the spectrogram of Figure 4.25 was used. In this figure multiple calls that overlap were identified.

The tracking of multiple pitches solved one problem, but posed another. The second pitch information array contained weaker pitch information that would normally be

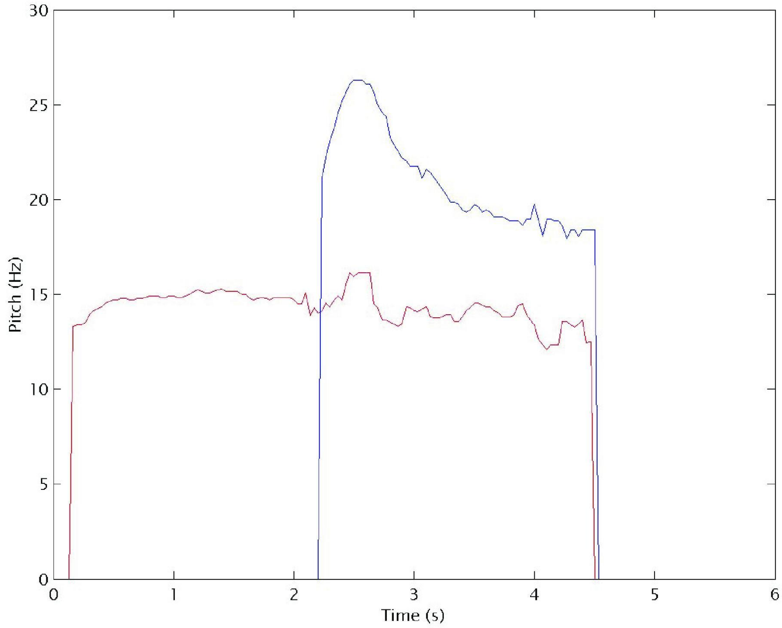

Figure 4.29: Dual pitch track algorithm output for the sound segment shown in Figure 4.25.

ignored, so that unwanted sounds that would not be detected by the single-track algorithm could be detected more easily when using the dual-track algorithm. The dual-track algorithm should generate a far greater number of false alarms.

4.3.6 Conclusion

In summary, the proposed algorithm will work reliably in some conditions, but not in all. Weak calls might not be detected if their harmonic structure is lost, and false alarms will increase when the dual-track algorithm option is selected. Either the single or the double pitch track algorithm can be selected depending on the need of the user.

The rumble detection algorithm identified the positions of rumbles in recordings with common background noises with an accuracy of 90.47%. It should, however, be remembered that both the manual detection of elephant rumbles and the definition of average data are subjective procedures which means that a good idea of the value of this algorithm could only be established by feedback from a number of different elephant researchers over a substantial period of time. The ability of the algorithm to determine the pitch of a rumble cannot be done accurately by inspection and could be a useful attribute.

The elephant recording collar that was developed in this study should solve the above mentioned problems. The collar records vocalizations very close to its source and this ensures that the harmonic nature of rumbles will not get lost. The problem of car engine and aeroplane noise should be solved to a large extent since the field of recording is now much smaller, meaning far-off car noises will have much less effect on the recording.

4.4 SUMMARY

Chapter 4 discussed the general experiments that were done with the recording collar. The first experiment was done to attain the best position for placement of the microphone. A number of different configurations were investigated by measuring the attenuation of sound over the required band of frequencies. The results of these experiments illustrated that one method provided much better results then the rest. An experiment was done to measure the power consumption of the collar and to test if the device remained stable at the extremes of the temperature range in which it was specified to operate. The results of the experiment showed that the device remained stable at the temperature at which it was tested and that 20 mA of electric current was drawn by the device.

The procedure of the field tests was explained and the results of the first field test were presented. Both the pleasing and the unsatisfying results were discussed. The recording process in the first field test ended prematurely because the battery ran down faster than had been expected due to a higher than anticipated power consumption of the microdrive that was used. The low battery power destabilized the reference voltage of the ADC causing a clicking sound at the beginning of every data buffer transfer. The second field test was done with new batteries and the microdrive was replaced with a CF card which solved the battery induced problem of the first field test. A number of elephant vocalizations were recorded during both tests. An unresolved problem was the difference in sensitivity between different recordings which was probably caused by blockages of the microphone at certain stages of the experiment.

The performance of the elephant rumble detection algorithm was tested by processing sounds with a different number of upper harmonics in varying sound levels. The results showed that the algorithm could detect sounds with three or more upper harmonics if the SNR was -8 dB or better. The detection algorithm was then tested on a number of elephant recordings and the results were compared to the manual analysis of the same recordings. The results showed that the algorithm detected 90.47% of the manually identified calls. It was also shown that the algorithm failed to detect rumbles that had lost their harmonic structure while noise sources with strong upper harmonics could cause false alarms. It was shown that the algorithm could be applied in a slightly different manner so that overlapping rumbles could be detected.