Defence Research and Development Canada
Recherche et développement pour la défense Canada

Evaluation of a Helmet-Mounted Auditory System for Gunfire Localization

Craig Burrell

Doug Saunders

Defence R&D Canada
Technical Memorandum
DRDC Toronto TM 2011-062
December 2011

Evaluation of a Helmet–Mounted Auditory System for Gunfire Localization

Craig Burrell
DRDC Toronto

Doug Saunders
DRDC Toronto

Defence R&D Canada – Toronto
Technical Memorandum
DRDC Toronto TM 2011-062
December 2011

Principal Author

Original signed by Dr. Craig Burrell

Dr. Craig Burrell
Defence Scientist, Individual Behaviour and Performance Section

Approved by

Original signed by Mr. Stephen Boyne

Mr. Stephen Boyne
Head, Individual Behaviour and Performance Section

Approved for release by

Original signed by Dr. Stergios Stergiopoulos

Dr. Stergios Stergiopoulos
Acting Chair, Knowledge and Information Management Committee
Acting Chief Scientist

© Her Majesty the Queen in Right of Canada as represented by the Minister of National Defence, 2011

© Sa Majesté la Reine (en droit du Canada), telle que représentée par le ministre de la Défense nationale, 2011

Table of contents

List of figures

List of tables

Acknowledgements

The authors would like to thank the following DRDC Toronto personnel: Dr. Stergios Stergiopoulos for initiating the research program on the Localization Array Processor, Mr. John Bowen for technical assistance in the laboratory, and Dr. Sharon Abel for providing comments on an earlier version of the manuscript. Mr. Andreas Freibert, of Advanced Ultrasound Technologies, also read and provided valuable feedback on the manuscript.

Abstract

The Localization Array Processor (LAP) is a system for acoustic localization of sources of rapid onset, short duration impulsive sounds, such as result from gunfire or explosions. The auditory environment is monitored by an array of helmet–mounted microphones. We report the results of a series of laboratory experiments to assess the performance characteristics of a prototype version of the LAP.

Three tests of the system were carried out. First, its operation in backgrounds of stationary noise was studied. It was found that the system's performance degraded rapidly as background noise levels increased, such that moderate levels of noise prevented the system from consistently identifying impulsive sounds in the environment. Second, tests of the direction–finding capability showed that the system was able to correctly identify the direction to the sound source in nearly 50% of cases, but that the angular resolution was not competitive with similar systems described in the literature. Finally, the recovery time of the system in a rapid fire scenario was studied, and it was found that it could detect an impulsive sound and refresh to a listening state in a few hundred milliseconds.

In general, the system's performance was satisfactory. Under optimal conditions it succeeded in identifying the azimuthal direction of origin of impulsive sounds, though the angular resolution provided by the current direction–finding algorithm is too coarse. The performance degraded in the presence of background noise. Numerous recommendations are made to guide future work in this area.

Résumé

Le processeur vectoriel de localisation est un système de localisation acoustique des sources de sons impulsifs de courte durée à apparition rapide, comme des coups de feu ou des explosions. L'environnement auditif est surveillé par un réseau de microphones montés sur casque. On présente les résultats d'un ensemble d'expériences de laboratoire pour évaluer les caractéristiques des performances d'une version prototype du processeur vectoriel de localisation.

On a réalisé trois essais sur le système. On a d'abord étudié le fonctionnement du système en présence de bruit stationnaire. On a découvert que les performances du système se dégradent rapidement avec l'augmentation de l'intensité du bruit de fond, et ce, à un point tel que des bruits d'intensité moyenne empêchent le système de repérer invariablement les sons impulsifs dans l'environnement. Ensuite, des essais sur la capacité de radiogoniométrie ont montré que le système peut repérer correctement la direction de la source des sons dans près de 50% des cas, mais que le pouvoir séparateur angulaire n'est pas aussi bon que celui de systèmes similaires décrits dans

la littérature. Enfin, on a étudié le temps de rétablissement du système dans un scénario de tir rapide, et on a trouvé que le système peut détecter un son impulsif et reprendre un état d'écoute en quelques centaines de millisecondes.

En général, les performances du système ont été satisfaisantes. Dans des conditions optimales, le système a réussi à déterminer l'azimut d'origine des sons impulsifs, mais le pouvoir séparateur angulaire fourni par l'algorithme de radiogoniométrie actuel est trop approximatif. Les performances se sont en outre dégradées en présence de bruit de fond. Enfin, un grand nombre de recommandations ont été formulées pour orienter des recherches futures dans ce domaine.

Executive summary

Evaluation of a Helmet–Mounted Auditory System for Gunfire Localization

Craig Burrell, Doug Saunders; DRDC Toronto TM 2011-062; Defence R&D Canada – Toronto; December 2011.

Background: The Localization Array Processor (LAP) is a system for acoustic localization of sources of rapid onset, short duration impulsive sounds, such as gunfire or explosions. The auditory environment is monitored by an array of helmet–mounted microphones. The LAP is designed to detect the acoustic muzzle blast produced by a firearm and identify the direction in the azimuthal plane from which it originated.

Principal results: We report the results of a series of laboratory experiments to assess the performance characteristics of a LAP prototype. In particular, three aspects of the system were studied: performance in the presence of background noise, accuracy of direction finding, and recovery time when exposed to rapid fire stimuli. For direction finding, it was found that in approximately 50% of cases the direction of origin of the impulsive sound was correctly identified as being within a 22.5° angular sector. This angular precision, however, is not competitive with that provided by similar systems reported in the literature, principally because of limitations of the direction–finding algorithm that is currently implemented. It was also found that the system could recover to a listening state within a few hundred milliseconds after detecting an impulse. The performance, however, degraded rapidly as the level of background noise increased; the device ceased to perform adequately in signal to noise ratios as high as 10 dB.

Significance of results: The results suggest that, though a credible technology, numerous improvements to the LAP are necessary to make it competitive with the state of the art. Improving the performance of the system in noisy environments will be necessary for the project's viability, as will improvements to the angular resolution. In general, the performance characteristics provide a reasonable but cautious warrant for future development of this technology.

Future work: Improvements to both the software and hardware would be necessary before the system could be ready for field use. The chief objectives of further work should be to improve the system's angular resolution, add capability to detect the shockwave produced by a supersonic bullet, improve performance in urban environments where reflections and echoes occur, improve performance in high noise environments, and reduce power consumption.

Sommaire

Evaluation of a Helmet–Mounted Auditory System for Gunfire Localization

Craig Burrell, Doug Saunders ; DRDC Toronto TM 2011-062 ; R & D pour la défense Canada – Toronto ; décembre 2011.

Contexte : Le processeur vectoriel de localisation est un système de localisation acoustique des sources de sons impulsifs de courte durée à apparition rapide, comme des coups de feu ou des explosions. L'environnement auditif est surveillé par un réseau de microphones montés sur casque. Le processeur vectoriel de localisation est conçu pour détecter la détonation acoustique d'une arme à feu et déterminer la direction de sa provenance dans le plan azimutal.

Résultats principaux : Nous présentons les résultats d'une série d'expériences en laboratoire pour évaluer les caractéristiques des performances d'un prototype de processeur vectoriel de localisation. Trois aspects du système sont notamment étudiés : les performances en présence de bruit de fond, la précision de la radiogoniométrie et le temps de rétablissement après une exposition à des stimuli de coups de feu rapides. On a déterminé que la précision de la radiogoniométrie était moyenne : dans environ 50% des cas, la bonne direction a été déterminée dans les limites de la précision maximale du dispositif. Par contre, la précision angulaire maximale ne concurrence pas celle des systèmes similaires décrits dans la littérature, surtout en raison des limites de l'algorithme de radiogoniométrie actuellement mis en place. De plus, on a trouvé que le système pouvait reprendre un état d'écoute en quelques centaines de millisecondes après la détection d'une impulsion. Les performances se sont toutefois dégradées rapidement avec l'augmentation de l'intensité du bruit de fond.

Après la présentation des résultats détaillés des expériences en laboratoire, on compare le processeur vectoriel de localisation à des systèmes similaires offerts dans le commerce ou décrits dans les ouvrages scientifiques. On formule ensuite diverses recommandations relatives à d'autres recherches sur cette technologie.

Portée des résultats : Selon les résultats, le processeur vectoriel de localisation est une technologie crédible pour localiser la direction d'origine d'événements acoustiques impulsifs, comme des coups de feu ou des explosions, mais il faut améliorer beaucoup cette technologie pour qu'elle puisse concurrencer les dispositifs de pointe. Pour que le projet soit viable, on devra améliorer les performances du système dans des environnements bruyants, ainsi que le pouvoir séparateur angulaire. En général, les caractéristiques des performances justifient de façon raisonnable le développement futur de cette technologie.

Recherches futures : Il faudrait améliorer le logiciel et le matériel pour que le système puisse être utilisé sur le terrain. Les principaux objectifs de recherches futures devraient être l'amélioration du pouvoir séparateur angulaire du système, l'ajout de capacités de détection de l'onde de choc produite par une balle supersonique, l'amélioration des performances dans des environnements urbains où se produisent des réflexions et des échos, l'amélioration des performances dans des environnements très bruyants, la réduction de la consommation d'énergie, la fourniture d'une alimentation portative, l'amélioration et la fixation sécuritaire du support des capteurs acoustiques, ainsi que la modification de l'interface utilisateur.

This page intentionally left blank.

1 Introduction

Common military equipment, such as helmets, communication headsets, and hearing protection, is known to impair the wearer's ability to accurately localize sound sources [1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12]. When, in a combat situation, a soldier hears gunfire or explosions, it is desirable that he or she ascertain as clearly as possible the direction from which those sounds originate.

Localization of the source of gunfire is a problem that has attracted considerable attention in both military and law enforcement contexts. Several different modalities have been proposed, including laser detection of rifle scopes [13, 14] and thermal analysis of muzzle flashes [15] or of bullets in flight [16]. The most common approaches, however, have involved acoustic measurements. Several systems using fixed [17, 18, 19] or vehicle–mounted [20, 21, 22] acoustic sensors have been described.

Soldier–worn acoustic sensors can also be deployed for gunfire localization, and several systems of this type have been reported. One system, developed by QinetiQ North America (McLean, Virginia) and called the Shoulder-Worn Acoustic Targeting System (SWATS) [23], consists of an acoustic sensor unit that is secured to the wearer's shoulder and is accompanied by an earpiece and small display. The SWATS system provides bearing and range to the shooter and has been recently deployed in both Iraq and Afghanistan [24]. A very similar system is Boomerang Warrior–X, from BBN Technologies (Cambridge, Massachusetts), which is also shoulder–mounted and is equipped with a comparable user interface [25].

Other systems embed multiple soldier–worn sensors within a wireless communication network for improved performance. A system from BBN Technologies [18] equipped each soldier with an array of twelve helmet–mounted acoustic sensors integrated with position– and orientation–tracking capabilities and connected by radio links to provide shooter localization for a group of mobile soldiers. More recently, a similar but more refined system, developed with funding support from Defense Advanced Research Projects Agency (DARPA), has been reported [26] that claims success with localization even in urban environments where acoustic reflections often confound localization efforts.

Defence Research and Development Canada (DRDC) Toronto has developed a prototype helmet–mounted microphone array for acoustic localization of sources of rapid onset, short duration impulsive sounds, such as small–arms muzzle blasts or explosions [27]. The resulting prototype system, called the Localization Array Processor (LAP), is designed to detect the acoustic signal produced by an impulsive source and identify the direction in the azimuthal plane from which it originated.

The purpose of this document is to report on the results of various performance tests carried out on the LAP prototype, to compare the performance to similar existing systems, and to make recommendations for future development of this technology.

2 Localization Array Processor (LAP)

The LAP prototype is shown in **Figure 1**. It is an integrated system, but is conceptually divisible into several components: (a) an input sensor array, (b) a signal processing unit, (c) a user interface, and (d) a power supply.

Figure 1: *The Localization Array Processor (LAP) prototype, showing the helmet–mounted microphone array, the signal processing unit, and the output display.*

The input sensor array consists of a roughly circular ring of sixteen microphones arrayed around the perimeter of an Army helmet. In the current prototype, the sensors are affixed to the helmet using Velcro pads, and face radially outward from the helmet's surface. Each sensor is separately wired to the signal processing unit. The microphones used in the current prototype have a sensitivity of −55 Decibels (dB) in the frequency range 100 Hz – 4 kHz, and can record sounds of up to 140 dB Sound Pressure Level (SPL) without clipping [27].

The signal processing unit consists of two parts: an acquisition unit and a processing unit. In the acquisition unit, the sixteen input channels are amplified, filtered, and input to two 24–bit, 8–channel Analog-to-Digital Converters (ADCs). The sampling rate is about 97.6 kHz, and all channels are sampled simultaneously. The resulting data stream, amounting to approximately 6 megabytes/s, is then passed to the processing unit.

The processing unit consists of a Digital Signal Processor (DSP) which monitors the input for signal peaks that could signify an impulsive acoustic wave. It is also capable of performing spectral analysis of the signal, although this is not currently used. When a candidate signal is detected, a direction–finding algorithm attempts to identify its direction of origin. The algorithm carries out cross–correlations between the channel in which the candidate was first detected and its neighbouring channels to determine the best choice for direction of origin. The result of this algorithm is then passed to the user interface.

The user interface is a Complex Programmable Logic Device (CPLD) that is flexible enough to supply output in a variety of different forms. In the present implementation it drives an array of sixteen Light Emitting Diode (LED) lights which are mounted, for convenience, on the processing unit. Each LED is associated with one of the input sensors; when an LED is illuminated, it indicates that an impulsive signal was detected and that its direction of origin is calculated to be in the angular sector associated with the corresponding sensor.

The entire system is powered by a single DC power supply. According to the contractors [27], it can operate with input voltages between $6 - 20$ V. In the laboratory, we used a voltage of 10 V.

With sixteen sensors, the best–case angular resolution of the system using the current direction–finding algorithm is $360\,^{\circ}/16 = 22.5\,^{\circ}$. The system assumes that the sound source is far enough away that a plane wave approximation is justified, and it also assumes that the impulsive acoustic wave has travelled directly to the helmet from the source. As such, it will be most accurate in relatively open spaces, where reflections from cliffs or buildings are not significant.

The system is designed to adjust its sensitivity based on the prevailing levels of stationary or quasi–stationary background noise: in a quieter background, it will respond to a lower amplitude signal. This feature, called variable gain control, will be discussed in more detail below (Section **3.2.3**) in the context of an experiment to which it is relevant.

In what follows, a stimulus/response model is used to describe the operation of the LAP. The *stimulus* refers to the acoustic signal presented, from a particular direction, to the input sensor array. The *response* refers to the pattern of illumination of the output LED array. For convenience, we sometimes say that “a sensor responds” to a stimulus; this means that the element of the LED array associated with that sensor was illuminated in response to the stimulus.

3 Experimental evaluation of LAP

In this Section, we describe laboratory experiments carried out on the LAP prototype system. The purpose of these experiments was to assess the performance of the system under controlled conditions. In particular, three aspects of the system were studied: performance in the presence of background noise, accuracy of direction finding, and recovery time when exposed to rapid fire stimuli.

3.1 Experimental apparatus

The experiments were all carried out in the Acoustics Workroom at DRDC Toronto. The basic arrangement is shown in **Figure 2**. The helmet–mounted sensor array was put on a dummy head and placed at the center of a circle (radius 1.5 m). Loudspeakers were placed around the perimeter of the circle, mounted at the same height as the helmet array. Although the loudspeakers could occupy sixteen different locations around the circle, in a one–to–one correspondence with the sensors on the helmet, in these experiments a maximum of five loudspeakers was used at any one time. The specific placement of the loudspeakers varied from experiment to experiment, and will be described in context below. The ambient sound level in the room was measured to be 75 dB SPL (52 dBA). This quasi–stationary ambient noise arose primarily from a mechanical room adjacent to the laboratory space.

The loudspeakers were driven by a workstation equipped with a multi–channel sound card (M–Audio (Irvine, California) Delta 1010LT [28]). Multi–channel audio files were prepared in pre–processing to generate the combination of signals required for testing (see **Annex A** for details), and were played using Windows Media Player. The signals output from the workstation passed through an eight–channel amplifier (BiAmp (Beaverton, Oregon) MCA 8150 [29]) before being sent to the loudspeakers. The LAP itself was powered by a Hewlett–Packard (Palo Alto, California) 6286A DC Power Supply.

3.2 Experiments

In this Section, we describe the particular experiments that were carried out and the results that were obtained.

3.2.1 Localization with background noise

The purpose of this experiment was to study the behaviour of the system as the Signal to Noise Ratio (SNR) changed. It was anticipated that the system would perform best at high SNR, and that performance would gradually degrade as the

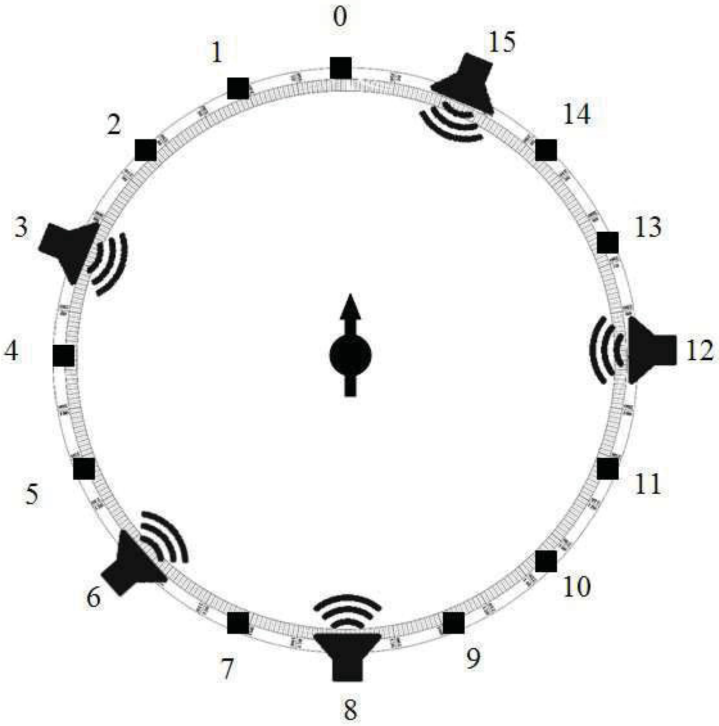

Figure 2: *Representative loudspeaker arrangement for LAP testing. The helmet–mounted sensor array was placed at the center of a circle (radius* 1.5 *m), with the helmet's forward direction indicated by the arrow. Loudspeakers could be positioned at sixteen evenly–spaced locations, labelled* $0-15$*, around the perimeter of the circle. Up to five loudspeakers were used at any one time. The specific placement of the loudspeakers varied from one experiment to another, and is described in context in each case.*

noise level increased. One objective of the experiment was to ascertain how quickly that degradation occurred. At some point, the noise was expected to be too loud for the impulsive stimuli to be detected.

For this experiment, four loudspeakers were used, in the arrangement shown in **Figure 3**. The azimuthal plane was divided into four contiguous zones, each associated with one of the four loudspeakers. The sensor array follows the contours of the helmet and is not exactly circular; it has a broader profile from the rear than from the front. In consequence, the zones have a slight forward – backward asymmetry. The zones define groupings of array sensors, as shown in **Table 1**.

A stimulus was played through one loudspeaker while the others simultaneously played a white noise background. The volume level of the white noise was adjusted so that the SNR at the location of the helmet had a particular value in the range -10 dB to 40 dB. The loudspeaker presenting the stimulus varied randomly; each

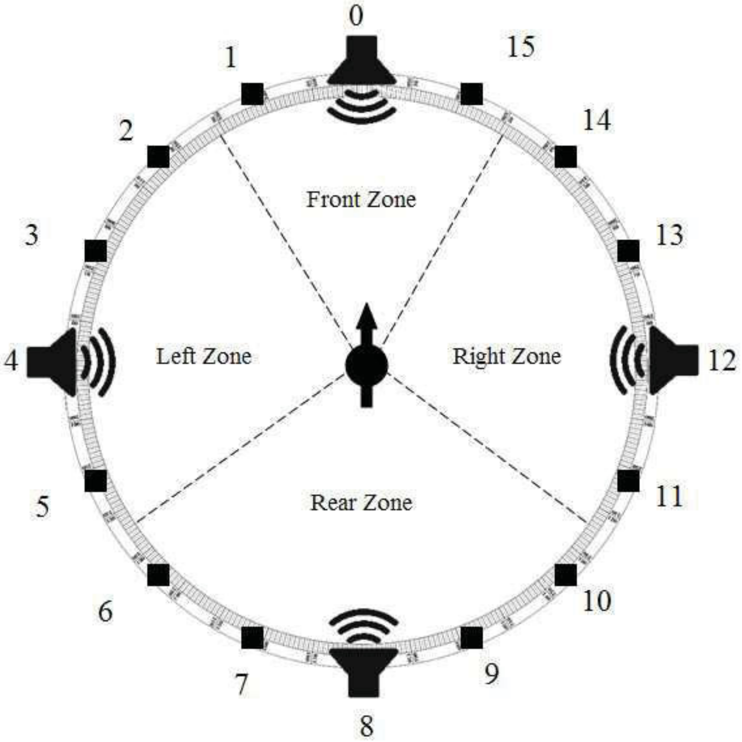

Figure 3: *Loudspeaker arrangement for testing localization in background noise. The arrow indicates the helmet's forward direction. The azimuthal plane is divided into four zones, one for each loudspeaker; the forward – backward asymmetry in the zones reflects the asymmetric shape of the helmet on which the sensors are mounted.*

loudspeaker played the stimulus twenty–five times at a given SNR. The stimuli were separated in time by at least 2 s, allowing ample time for the LAP to process the previous stimulus (see Section **3.2.3**) and for the response to be recorded.

Two different impulsive stimuli were used: the first a gunshot recording with a short, sharp profile (shown in **Figure 4**, left hand side), and the second a more diffuse explosive sound (shown in **Figure 4**, right hand side). The peak amplitude of Stimulus 1 at the location of the sensors was measured to be 91 ± 1 dB SPL (90 ± 1 dBA); the corresponding value for Stimulus 2 was found to be 94 ± 1 dB SPL (91 ± 1 dBA).

If the stimulus was presented by a loudspeaker and exactly one of the sensors in the associated zone responded, it was counted as a "correct" response. If exactly one sensor responded but it was not in the correct zone, it was counted as an "incorrect" response. If more than one sensor responded, it was counted as a "multiple" response, which we considered to be a type of incorrect response. If no sensor responded, it was counted as a "null" response.

Zone	Sensors
Front	0, 1, 15
Left	2, 3, 4, 5
Rear	6, 7, 8, 9, 10
Right	11, 12, 13, 14

Table 1: *Sensors associated with each zone for testing localization with background noise. The zones are illustrated in* **Figure 3**.

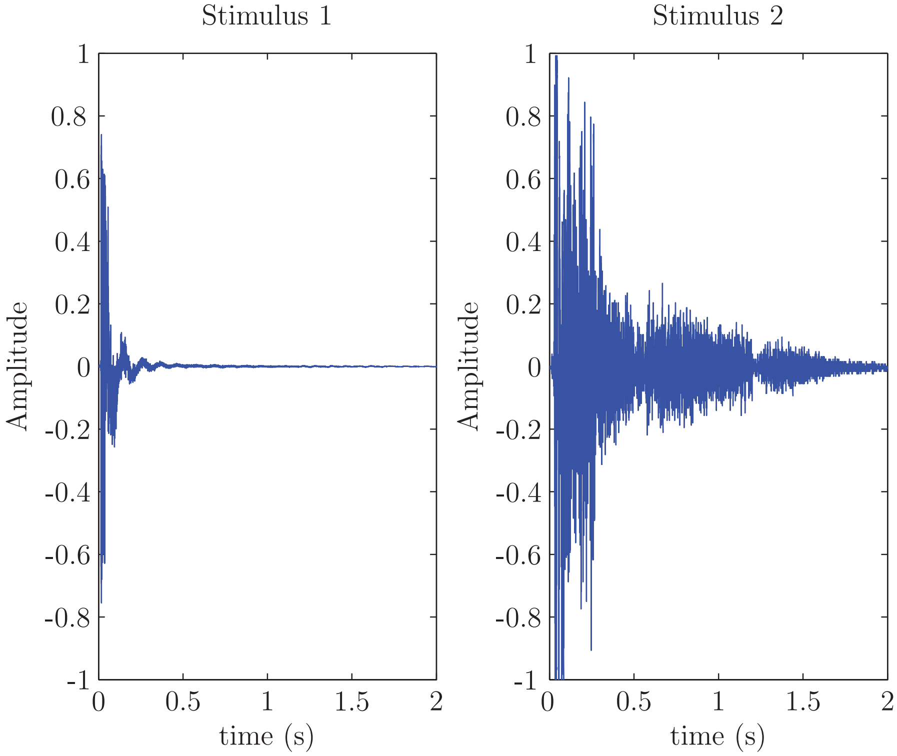

Figure 4: *Impulsive stimuli used to test the LAP.*

The results for Stimulus 1 (**Figure 4**, left hand side) are shown in **Figures 5** and **6**. The bar chart in **Figure 5** shows the prevalence of correct, incorrect, multiple, and null responses as a function of SNR. We see that the rate of correct response declined steadily in the range -10 dB $<$ SNR $<$ 0 dB, with a concomitant rise in the null response rate. The rate of incorrect responses was low across the SNR values that were tested. It is important to note that even at those low SNR values for which the LAP had ceased to respond, the unaided ear could easily identify the stimulus against the background noise.

The radar diagram in **Figure 6** takes the correct responses from **Figure 5** and illustrates their directional dependence, plotting the correct–response rate in each of

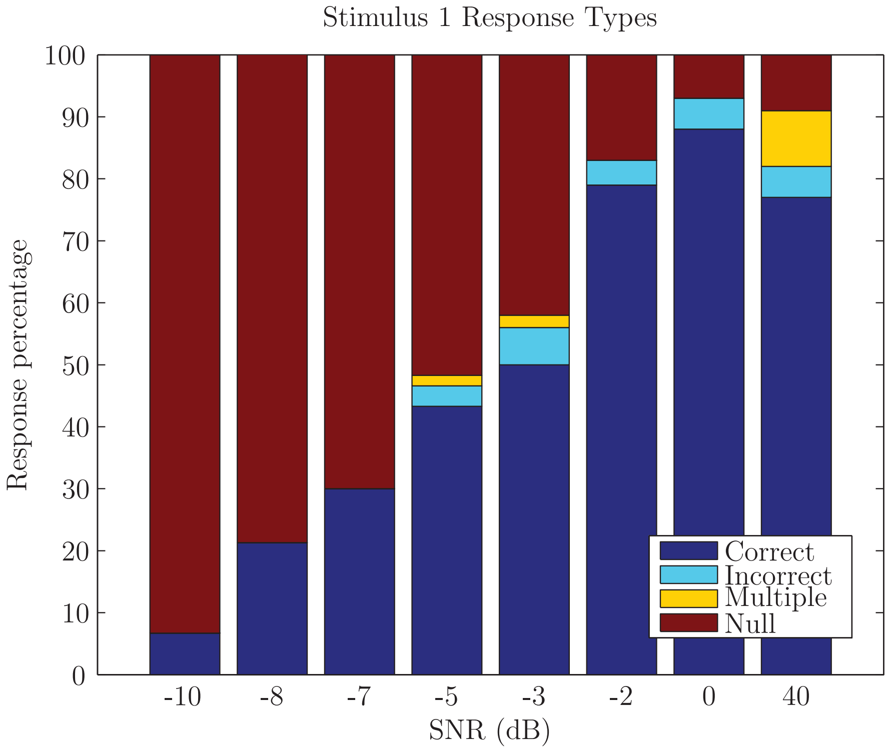

Figure 5: *Response types received for Stimulus 1 at various SNR values.*

the four zones for several SNRs. The results show that the correct–response rate was relatively high ($> 80\%$) when the SNR was high (0 dB), and declined at lower SNRs. With this stimulus, the response degraded rapidly between -3 dB $<$ SNR < -2 dB. The response profile of the helmet was approximately symmetric, though the right–hand side of the helmet array retained higher performance down to SNR $= -3$ dB than did the other sides. At lower SNR values, one or more of the correct–response rates was zero, and so data for these values are not shown in **Figure 6**. We also observed poor performance on the left–hand side of the helmet for the case when there was no white noise background added (SNR = 40 dB, the noise coming exclusively from ambient noise in the laboratory), and this was due to an unusually high null response rate. The reasons for this anomaly are not known, but are perhaps attributable to a behaviour that was occasionally observed during data collection, in which the LAP would cease to respond and have to be reset (see Section **4** below).

The results obtained for Stimulus 2 (**Figure 4**, right hand side) are shown in **Figures 7** and **8**. The data in **Figure 7** are broadly similar to what was observed for the previous stimulus in **Figure 5**. Notice, however, that for this stimulus, the overall decline in performance happens at higher SNR values than in the previous case: at and below SNR = 10 dB the LAP has ceased to respond. The reasons for this

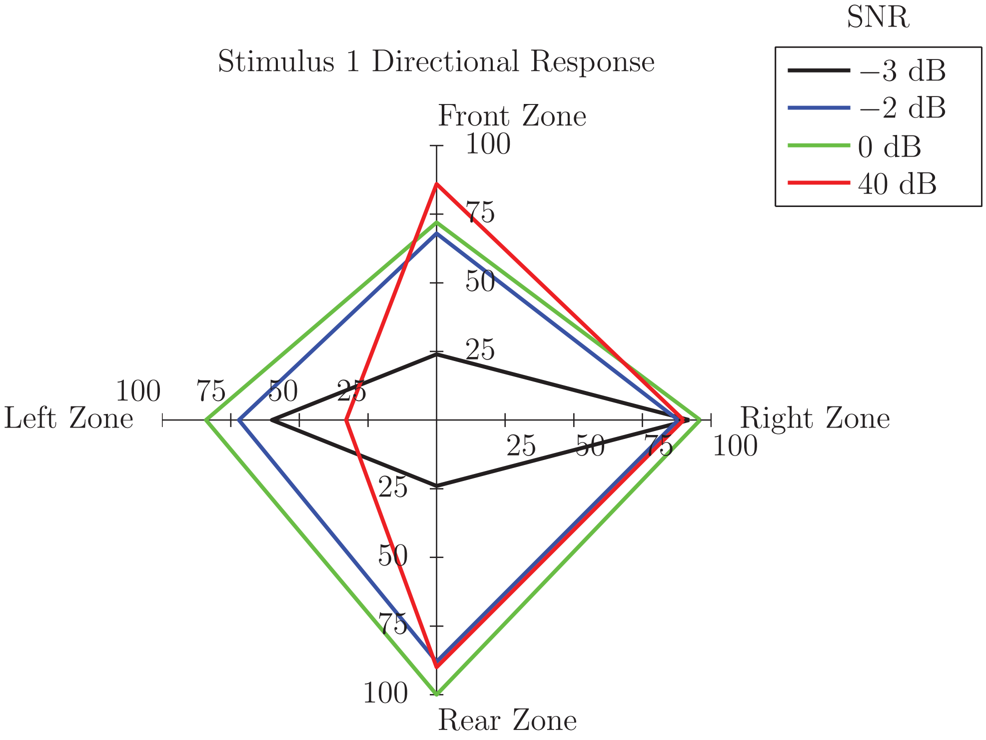

Figure 6: *Directional response of the LAP to Stimulus 1, showing the percentage of correct responses in each zone at various SNR values. Lower SNR values are omitted because the percentage of correct responses was zero in one or more zones.*

difference are not clearly understood, but are likely caused by the different temporal profiles of the two stimuli (**Figure 4**). The broader profile of Stimulus 2 appears to be a less effective trigger for the impulse detection algorithm. Multiple responses are a problem at very high SNR, but with relatively low levels of background noise the LAP correctly identifies the direction fairly reliably ($> 80\%$). At lower SNR, the null response rate increases rapidly. As above, we note that at these noise levels it is easy for the unaided ear to pick out the stimulus against the background noise. The rate of incorrect responses was higher for this stimulus than in the previous case (sometimes as high as 10%).

The radar diagram in **Figure 8** shows the directional dependence of the correct responses from **Figure 7**, plotting the correct–response rate in each of the four zones for several SNRs. As in the previous case, the right–hand side of the sensor array sustained good performance to lower SNR values than the other sides. The poor performance at SNR = 40 dB reflects the low rate of correct responses at that SNR value, as shown in **Figure 7**.

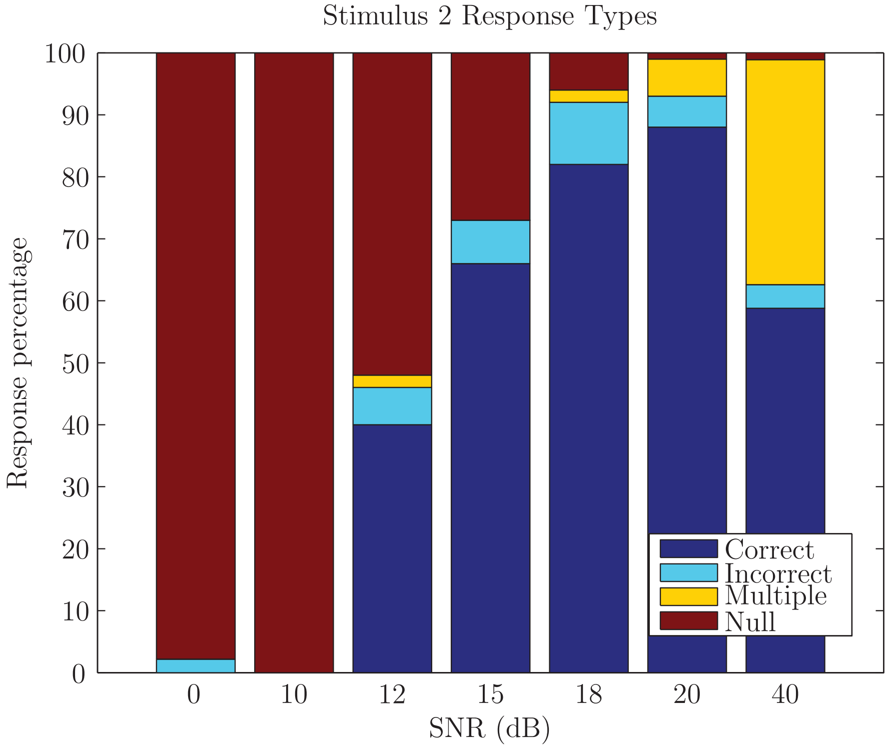

Figure 7: *Response types received for Stimulus 2 at various SNR values.*

The LAP sometimes identified more than one direction of origin in response to a single stimulus (thus creating the "Multiple" response type in **Figures 5** and **7**). These responses can be further subdivided based on how many of the multiple responses were in the correct zone, and these data are shown in **Table 2**. In approximately one–quarter of such cases, all of the responses were in the correct zone (see the first row in **Table 2**), but approximately three–quarters of the time at least one response was in the correct zone and at least one was in an incorrect zone (see the second row in **Table 2**). As the final row of the Table indicates, it was rare for all of the responses to indicate the wrong zone. (This occurred in just one case.) The reason

Distribution of Multiple Responses	Stimulus 1	Stimulus 2
Correct Zone Only	25%	27%
Correct and Incorrect Zones	67%	73%
Incorrect Zones Only	8%	0%

Table 2: *Distribution of LAP responses when multiple responses were received for a single stimulus.*

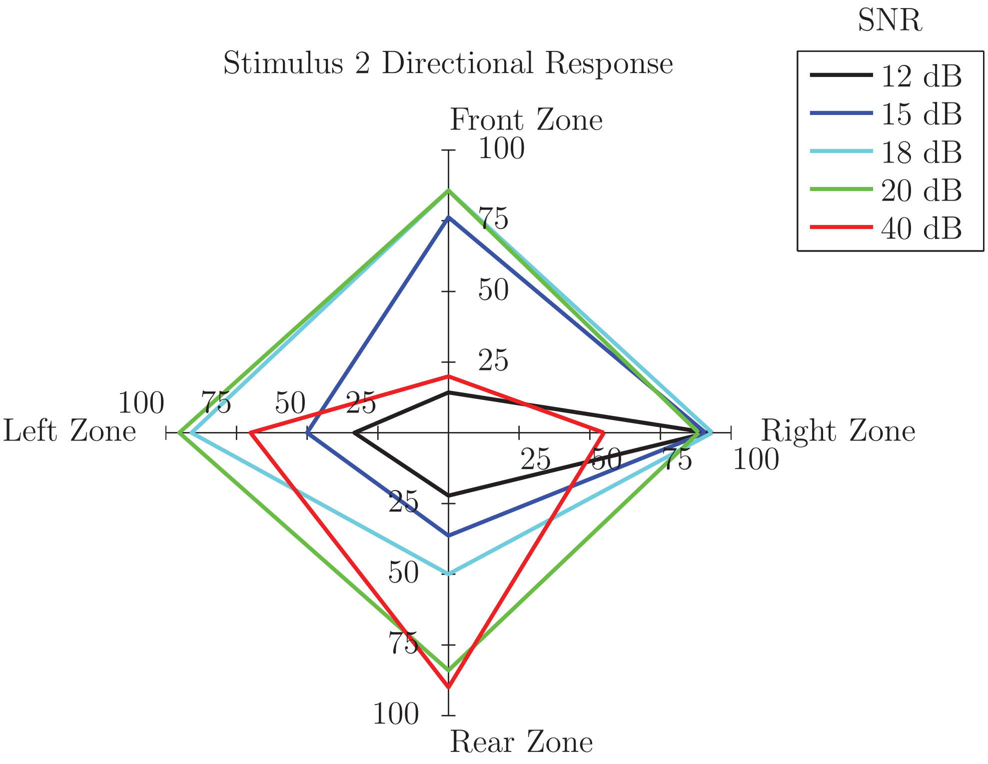

Figure 8: *Directional response of the LAP to Stimulus 2, showing the percentage of correct responses in each zone at various SNR values. Lower SNR values are omitted because the percentage of correct responses was zero in one or more zones.*

that the LAP sometimes indicates more than one direction of origin in response to a single stimulus is not certain. The fact that such responses occur most frequently at the lowest noise levels (see **Figures 5** and **7**) suggests that it may be triggering on low–level environmental sounds that would otherwise be masked by noise. It is even possible that the array is picking up reflections of the initial stimulus from the laboratory walls. Alternatively, there may be a problem with the signal processing algorithms when noise levels are low.

The most significant finding from this experiment is that the LAP does not respond well when background noise levels compete in intensity with the impulsive stimuli. We observed that the device ceased to detect the stimuli at noise levels for which the unaided ear could still hear them clearly. It may be possible to improve this aspect of the system's performance by making changes to the signal processing algorithms. Investigating this possibility should be a priority for future development.

3.2.2 Localization resolution

The purpose of this experiment was to study the accuracy with which the LAP was able to identify the direction of origin of an impulsive acoustic stimulus. The stimulus used in this experiment was Stimulus 1 from **Figure 4**; it was chosen instead of Stimulus 2 (also shown in **Figure 4**) because, as discussed in the previous Section, the LAP responded more reliably to it.

Figure 9: *Loudspeaker arrangement for testing localization resolution. The black arrow indicates the helmet's initial forward direction, and the grey arrows the directions to which the helmet was rotated as the experiment proceeded.*

The physical arrangement of loudspeakers used in this experiment is illustrated in **Figure 9**. Because not enough loudspeakers were available to occupy all sixteen locations simultaneously, segments of the sensor array were tested sequentially using an arc of five loudspeakers, as shown. The loudspeakers were placed at angular intervals of 22.5°, in correspondence with the angular spacing of the sensors in the array, such that each loudspeaker "targeted" its associated sensor. The helmet was rotated by 90° after each session (as indicated by the grey arrows in **Figure 9**) to allow for testing of the entire sensor array.

Impulsive stimuli were directed at the array at intervals of 4 s. Building on experience gained in the previous experiment, white noise was added to the background at a low level (SNR = 20 dB) to decrease the number of occasions on which multiple sensors responded to a single stimulus. A block of thirty stimuli was presented by each loudspeaker and the responses of the array were tabulated.

The results are shown in **Figures 10** and **11**. Each histogram shows how the LAP responded when a particular sensor (identified in the histogram's title) was targeted. The histograms show single–sensor responses (in blue), multiple–sensor responses (in green), and null responses (in red). Ideally, only single–sensor responses would be observed, with all of the responses confined to the target sensor's bin.

In practice, we observed that single–sensor responses, though focused around the target sensor, were distributed across other directions as well. Multiple–sensor responses and null responses were also observed, though in most cases these did not predominate. Some sensors exhibited behaviour that closely approximated the ideal (such as Sensors 8 and 15), while in one case (Sensor 11), the LAP failed to respond correctly even once. (We conclude that a hardware problem had disabled this sensor's signal.)

In several cases (Sensors 10 and 13) a high percentage of null responses was observed. The reason for this is not well understood, but may be related to the transient unresponsiveness of the array that was occasionally observed in the previous experiment (see also Section **4**).

In several cases, the target sensor responded rarely, and a neighbouring sensor responded instead. This is possibly due to a deficiency in the mounting of the sensors in the current prototype. Ideally, they would be mounted flush against the surface of the helmet so that they faced radially outward. In practice, however, the sensors are sometimes slightly askew and face a direction that only roughly corresponds to the radial directions (see **Figure 12**). Improving the mounting of the sensors would likely improve the angular resolution of the LAP, and should be considered a high priority in further development. (This observation was already made in the developer's Contract Report [27].)

To quantify the accuracy of the LAP's direction finding, we compute for each histogram in **Figures 10** and **11** a set of related performance measures that describe the degree to which the array's response is focused around the target sensor. We first define the quantity

$$\delta_i = \frac{1}{M - n_i} \sum_{j=0}^{N-1} n_j d(i,j), \tag{1}$$

where n_i is the number of responses recorded from sensor i, M is the total number of stimuli to which the array responded, N is the total number of sensors in the array ($N = 16$ in this case), and $d(i,j)$ gives the shortest distance, in number of hops, between the target sensor and the sensor that actually responded. (Because the array is circular, $d(i,j) \leq \lfloor N/2 \rfloor$.) The quantity δ_i is the average distance, in

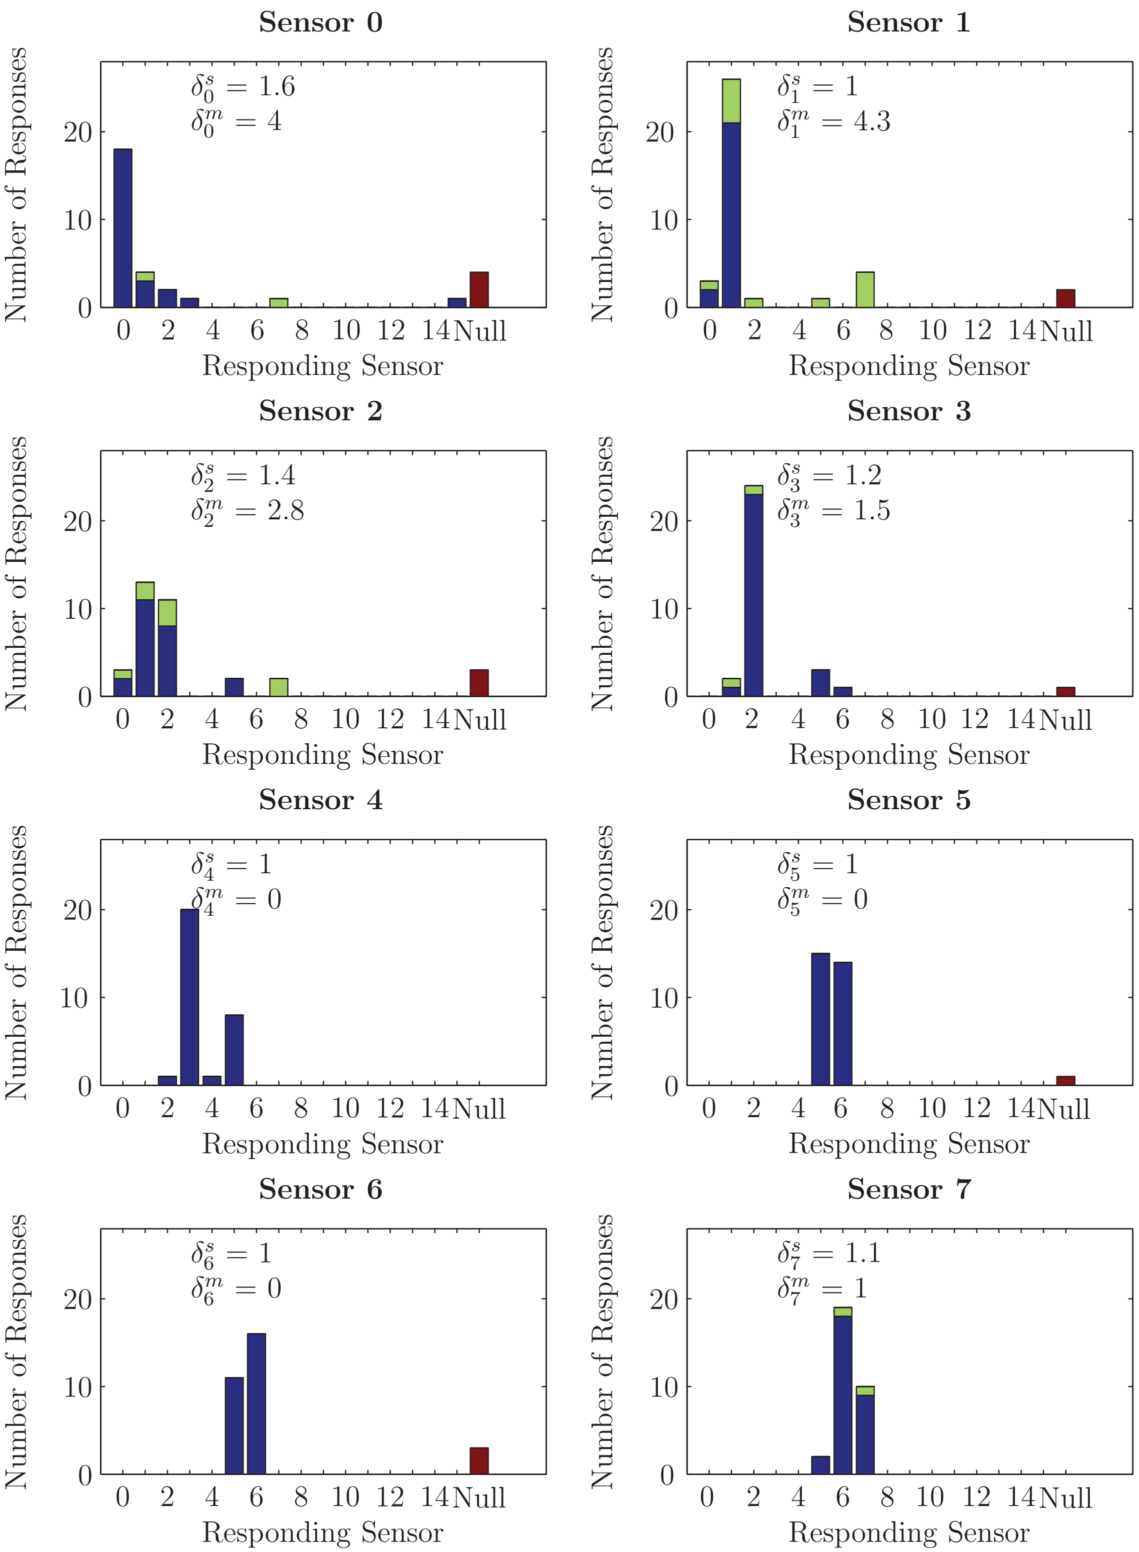

Figure 10: *Resolution test results for target Sensors 0-7. Single sensor responses are indicated in blue, and multiple responses in green. Null responses are indicated in red.*

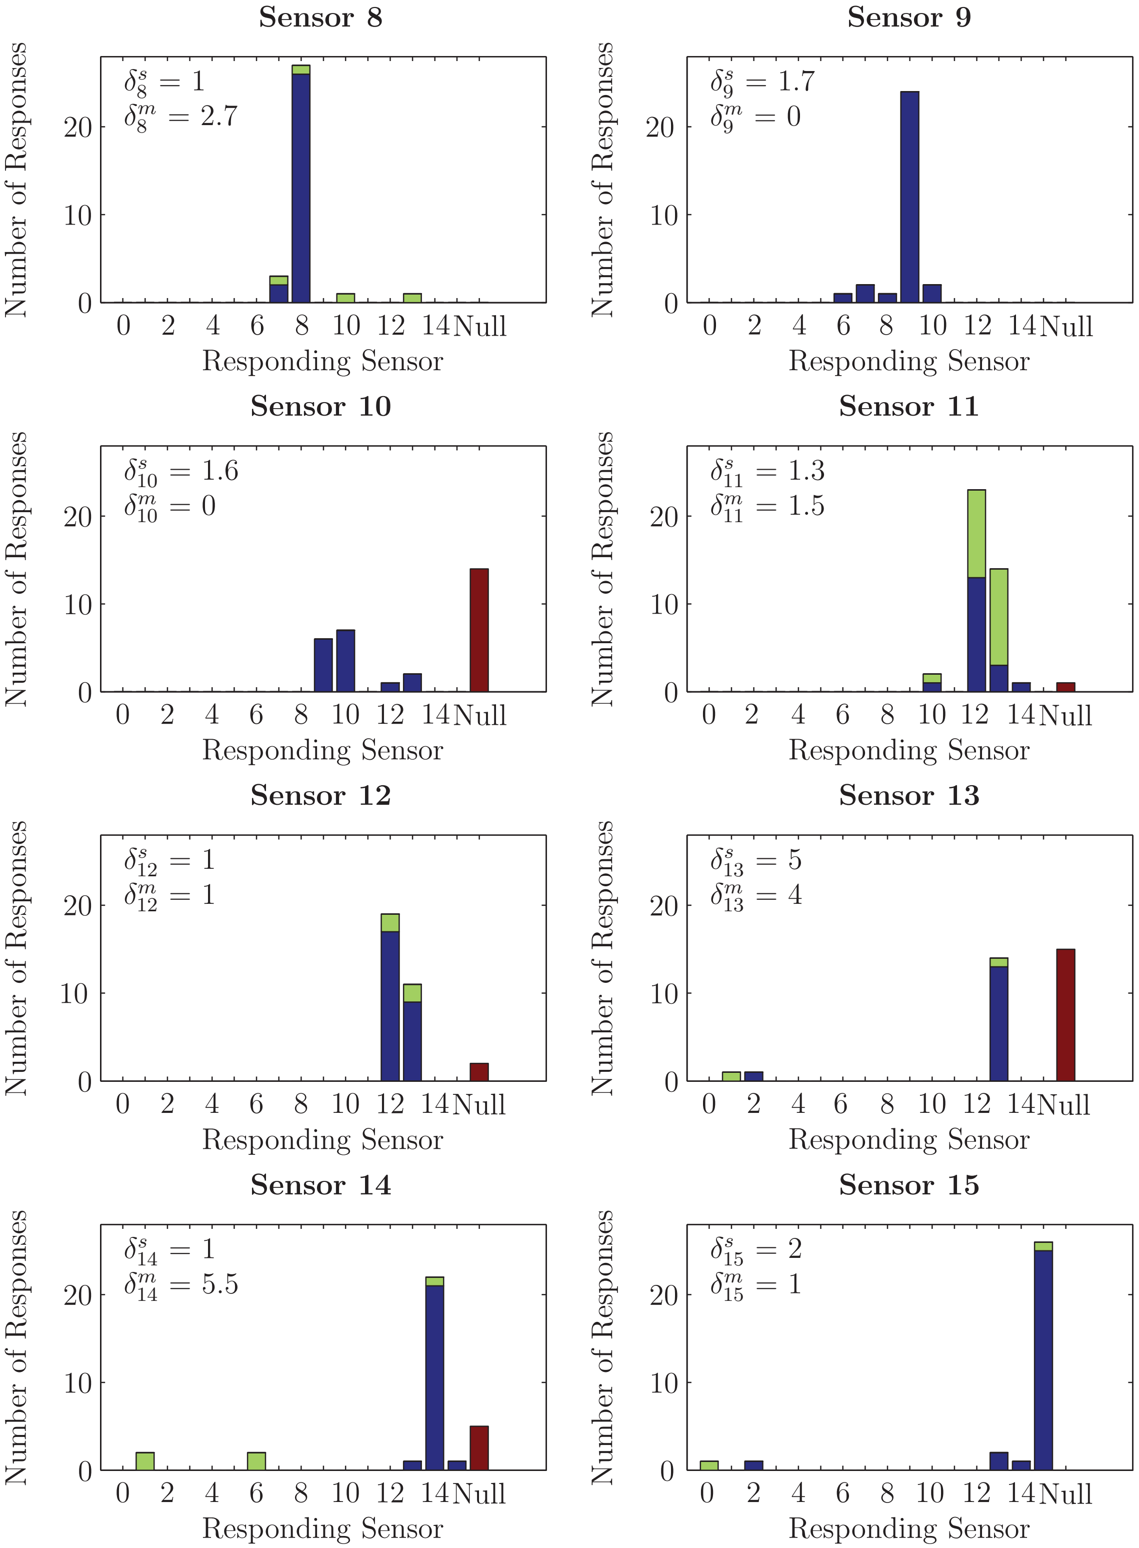

Figure 11: *Resolution test results for target Sensors 8-15. Single sensor responses are indicated in blue, and multiple responses in green. Null responses are indicated in red.*

Figure 12: *Detail of sensor mounting on LAP prototype. The microphone is not flat, resulting in skewed angular coverage.*

number of hops, from the target sensor to the responding sensor when the response was incorrect (hence the division by $M - n_i$, the number of incorrect responses). This quantity gives a measure of how unfocused were the incorrect responses.

A second performance measure folds in information about how frequently an incorrect response was given. We define a ratio $r_i = (M - n_i)/M$ for each sensor, giving the fraction of cases for which a sensor other than the target sensor responded. Then the quantity

$$\Delta_i = r_i \delta_i = \frac{1}{M} \sum_{j=0}^{N-1} n_j d(i,j) \tag{2}$$

gives the average deviation, in number of hops, from the target direction, taking into account both correct and incorrect responses. If all responses were correct, we would have $\Delta_i = 0$, and larger values of Δ_i indicate poorer accuracy in direction–finding.

We can convert these quantities into a measure of angular focus by defining

$$\Theta_i = \frac{360^\circ}{N} \left(1 + P(N)\Delta_i\right), \tag{3}$$

where

$$P(N) = \frac{N(N-1)}{\lfloor N/2 \rfloor (\lfloor N/2 \rfloor + (N \bmod 2))} = \begin{cases} \frac{4(N-1)}{N} & ; N \, \text{even} \\ \frac{4N(N-1)}{N^2-1} & ; N \, \text{odd} \end{cases} \tag{4}$$

in the general case of an N–element circular array. If the array always responded correctly, we would have $\Delta_i = 0$ and therefore $\Theta_i = 360\,^\circ/N$, indicating that the sensor covers the sector alloted to it and nothing more. If, on the other hand, we had a worst–case response profile such that, for a stimulus presented to sensor i, all of the sensors responded with equal frequency, we would have, from (2),

$$\begin{aligned}\Delta_i &= \frac{1}{M}\frac{M}{N}\sum_{k=0}^{N-1} d(i,k) \quad \text{since}\, n_k = \frac{M}{N}\,\forall\, k \\ &= \left\{\begin{array}{ll} \frac{N}{4} & ;N\,\text{even} \\ \frac{N^2-1}{4N} & ;N\,\text{odd}\end{array}\right. ,\end{aligned}$$

which, when combined with (3) and (4), yields

$$\Theta_i = \left\{\begin{array}{l} (N\,\text{even})\ \frac{360\,^\circ}{N}\left(1+\frac{4(N-1)}{N}\frac{N}{4}\right) \\ (N\,\text{odd})\ \frac{360\,^\circ}{N}\left(1+\frac{4N(N-1)}{N^2-1}\frac{N^2-1}{4N}\right)\end{array}\right\} = 360\,^\circ, \tag{5}$$

indicating that the array provides no useful angular discrimination in this case. (The form of $P(N)$ in (4) has been chosen so that this limit case gives this result.)

i	n_i	δ_i^s	Δ_i^s $(\times 10^{-2})$	Θ_i^s	δ_i^m
0	18	1.6	6.3	28 °	4.0
1	21	1.0	4.4	26 °	4.3
2	8	1.4	6.1	28 °	2.8
3	0	1.2	4.3	26 °	1.5
4	1	1.0	3.5	25 °	–
5	15	1.0	3.5	25 °	–
6	16	1.0	3.7	26 °	–
7	9	1.1	3.8	26 °	1.0
8	26	1.0	3.6	26 °	2.7
9	24	1.7	5.6	27 °	–
10	7	1.6	9.7	31 °	–
11	0	1.3	7.1	28 °	1.5
12	17	1.0	3.9	26 °	1.0
13	13	5.0	36	53 °	4.0
14	21	1.0	4.4	26 °	5.5
15	25	2.0	6.9	28 °	1.0

Table 3: *Performance measures for LAP direction finding, based on data shown in* **Figures 10** *and* **11**. *The index i labels the channel number. Other column headings are defined in the text.*

With these definitions in place, we can compute performance measures for the data shown in **Figures 10** and **11**. The results are given in **Table 3**. The second column gives the number of times, out of thirty, that the correct sensor (and only the correct sensor) responded to the stimulus; the average of these values is 14, indicating that the array responded correctly slightly less than half the time. If we also include single responses from sensors adjacent to the target sensor then the average rises to 23 (77%).

For the average displacement from the target direction when the result is incorrect, as defined in (1), we divide the data into two groups consistent with the division shown in **Figures 10** and **11**: single response cases, denoted by δ_i^s, and multiple response cases, denoted by δ_i^m. For the single response cases, the values of δ_i^s show that when the target direction is not correctly identified, the array nonetheless tends to respond with a direction close to the target, with the deviation being usually between 1 and 2 hops. In other words, the array fairly consistently identifies the correct zone (see **Figure 2**), even when the precise direction within that zone is not identified.

The small values of Δ_i^s shown in **Table 3** indicate that a direction close to the correct direction is usually identified. When these values are converted into an angular measure Θ_i^s according to (3), we see that, for the most part, the angular sector to which the sensor responds is comparable to the best–case result of $22.5\,^\circ$. The exception to this is Sensor 13, for which the incorrect responses are not clustered near the correct channel.

For those cases in which multiple sensors responded to a single stimulus, the data, shown in the final column of **Table 3**, indicate that the responses tend to be distributed broadly around the array, rather than focused near the target direction. The proportion of cases in which multiple responses occur is small, but that they occur at all is a problem that further developments of this technology should try to correct.

3.2.3 Localization in rapid fire

The purpose of this experiment was to test the response of the array to multiple stimuli separated by a short time interval in order to ascertain how long the system requires to return to the listening state following the detection of a stimulus.

The experimental setup was as shown in **Figure 13**. Two identical stimuli, with start times separated by an interval Δt, were presented to the array at right angles. Stimulus 1 (**Figure 4**, left hand side) was used, and white noise was added to achieve an overall SNR of 0 dB. This SNR was selected in order to maximize the array's response rate (see **Figure 5**). Several time intervals, Δt, in the range 25 ms to 1000 ms were tested; fifty stimuli pairs were presented for each time interval; the order in which the time intervals were tested did not affect the outcomes.

Figure 13: *Loudspeaker arrangement for testing rapid fire localization. The arrow indicates the helmet's forward direction.*

The results of the tests are shown in **Tables 4** and **5**. In these Tables, the response of the LAP to the two stimuli is signified by an ordered pair of boxes in which a check–marked box (☑) signifies a correct response, a crossed–out box (☒) signifies an incorrect response, and a blank box (☐) signifies a null response. A response was considered correct if a sensor in the zone of the array facing the stimulus direction responded, and incorrect if a sensor responded in some other zone of the array. The data indicate that incorrect responses were rare.

The two most common types of response were ☑☑ (correct responses to both stimuli) and ☑☐ (correct response to the first stimulus and a null response to the second). The relative prevalence of these two responses was strongly dependent upon the Δt separating the stimuli. The results for these response types, together with the double null (☐☐) responses, are shown in **Figure 14**. For relatively long intervals, the ☑☑ response dominates, indicating that the array is able to detect the first stimulus and recover to a listening status in order to detect the second stimulus with a fair degree of reliability. As the interval decreases, however, the array's ability to detect the second stimulus drops off. In a certain range ($\Delta t \sim 50 - 150$ ms), the second stimulus is almost never detected. The asymmetry between ☑☐ and ☐☑ responses indicates that the effect is not simply due to the fact that the array occasionally fails to respond to a stimulus; rather, the processing associated with the first stimulus is evidently interfering with the array's ability to detect the second stimulus. Interestingly, at

Response	Interval between stimuli, Δt (ms)													
	25	30	35	40	50	75	100	150	200	300	400	500	750	1000
☑☑	0	29	4	1	0	0	0	0	22	27	20	19	32	44
☑☒, ☒☑	2	0	2	1	0	0	0	0	0	3	3	3	9	1
☒☒	0	0	1	0	0	0	0	0	0	1	0	0	0	0
☑☐	40	20	39	43	45	42	45	45	22	16	20	21	5	2
☐☑	4	0	0	3	3	5	1	0	4	2	5	7	3	2
☒☐, ☐☒	0	0	0	0	0	0	0	0	0	0	0	0	0	0
☐☐	4	1	3	2	2	3	4	5	0	0	1	0	0	0
> 2 responses	0	0	1	0	0	0	0	0	2	1	1	0	1	1

Table 4: *Rapid fire test results: number of responses in each category as a function of time interval betweens stimuli. First shot: rear; second shot: left. The response categories are explained in the text. Fifty trials were run for each time interval.*

still shorter intervals ($\Delta t \sim 30 - 40$ ms), the ☑☑ response makes a partial comeback before disappearing again at $\Delta t = 25$ ms.

To understand these results, we must make reference to two technical features of the LAP: *deaf time* and *variable gain control*. The deaf time of the array is a period during which the array ceases to monitor the environment following the detection of a stimulus. If the system detects one stimulus, it will never detect another if it presents at an interval briefer than the deaf time. Deaf time is a tunable parameter which was set during these experiments to 30 ms. This explains why the system always failed to detect both stimuli when $\Delta t = 25$ ms.

Variable gain control is a feature of the LAP that allows it to adapt its sensitivity to the level of ambient noise [27]. The signals collected by the sensors are subjected to a gain change applied by the data acquisition component of the signal processing unit. When the ambient noise level is low, the gain applied to the input signal is high, making the system sensitive even to impulsive stimuli of low intensity. When the noise level increases, however, the gain automatically adjusts to a lower level in order to dampen the noise and prevent clipping. Of course, it also results in the system being less sensitive to incident stimuli.

These characteristics of the variable gain control system account for the general features of **Figure 14**. Upon detection of the first stimulus, the variable gain control system, responding to the increased signal power, reduces the array's sensitivity. When the first stimulus has passed, however, and the noise level has reverted to its earlier value, the sensitivity returns linearly to its previous state. The net result is a period of reduced sensitivity following the detection of an impulsive signal. Thus, referring to **Figure 14**, when the Δt is very short, the system detects both stimuli

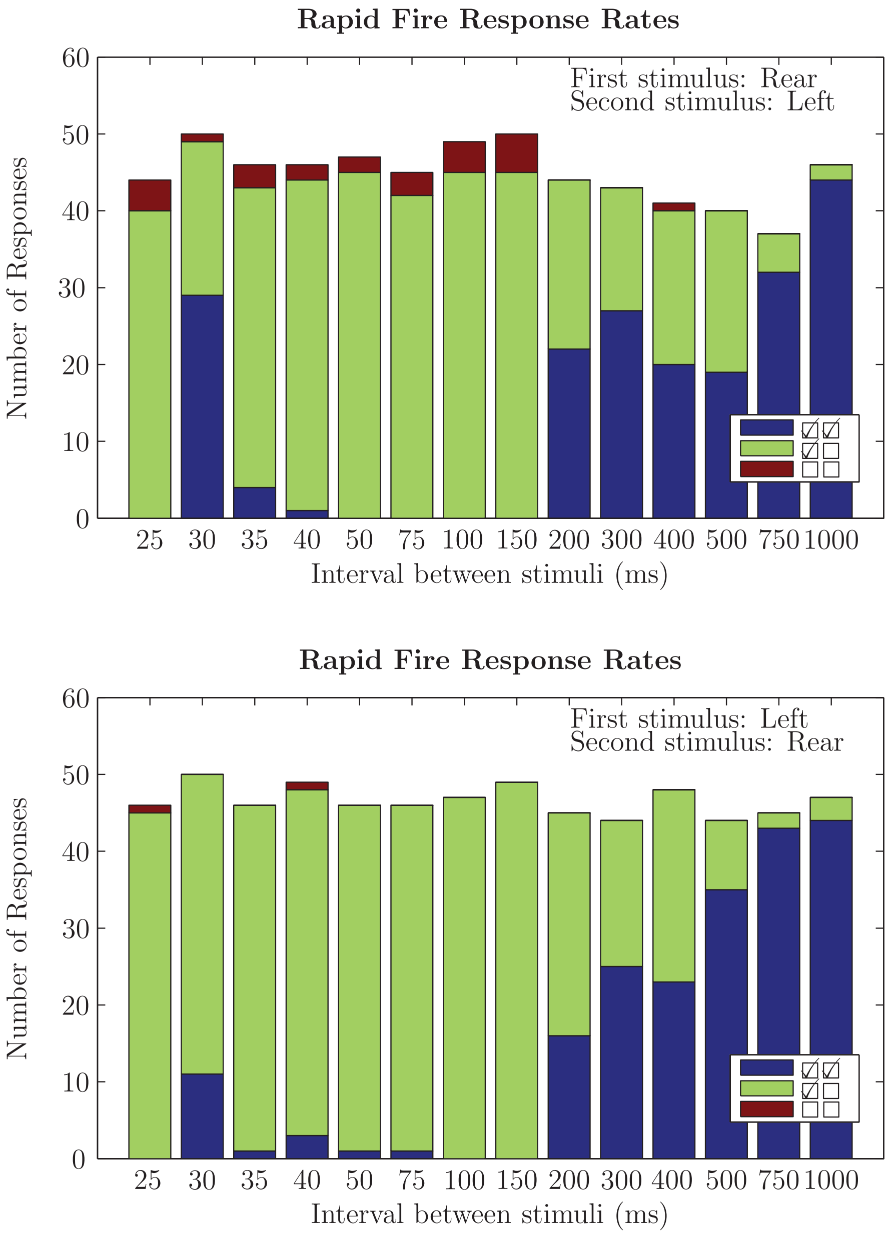

Figure 14: *Selected results of rapid fire tests. The top figure shows results when the first stimulus is from the rear and the second from the left-hand side. In the bottom figure the order of the stimuli is reversed. The results are described by an ordered pair of boxes: ☑☑ indicates that both stimuli were correctly detected, ☑☐ indicates that only the first was detected, and ☐☐ indicates that neither was detected. For a complete tabulation of results, see* **Tables 4** *and* **5**.

Response	Interval between stimuli (ms)													
	25	30	35	40	50	75	100	150	200	300	400	500	750	1000
☑☑	0	11	1	3	1	1	0	0	16	25	23	35	43	44
☑☒, ☒☑	0	0	0	0	0	0	0	0	2	3	0	5	5	3
☒☒	0	0	0	0	0	0	0	0	0	0	0	0	0	0
☑☐	45	39	45	45	45	45	47	49	28	19	25	9	2	3
☐☑	2	0	2	1	4	4	3	1	4	2	2	0	0	0
☒☐, ☐☒	2	0	2	0	0	0	0	0	0	0	0	0	0	0
☐☐	1	0	0	1	0	0	0	0	0	0	0	0	0	0
>2 responses	0	0	0	0	0	0	0	0	0	1	0	1	0	0

Table 5: *Rapid fire test results: number of responses in each category as a function of time interval betweens stimuli. First shot: left; second shot: rear. The response categories are explained in the text. Fifty trials were run for each time interval.*

before the sensitivity is substantially reduced, and when the Δt is long, the system detects both stimuli because the sensitivity has had time to return to its original value. For intermediate values, however, the second stimulus is not detected because, being the same amplitude as the first stimulus, it does not stand out above the "noise floor" established by the first stimulus.

The rapid fire experiment was repeated, reversing the order of the stimuli, to rule out a possible bias resulting from variable sensor sensitivity in different parts of the array. As is clear from **Figure 14**, substantially the same pattern of responses is found for both orderings.

These findings indicate that the LAP can detect two stimuli with a fair degree of reliability ($>50\%$) when the stimuli are separated by at least 0.5 s. When the separation grows to 1.0 s, the success rate is high ($>90\%$). These findings apply when conditions are near optimal; in the presence of increased noise, the rate of success would be reduced, as was found in Section **3.2.1**.

4 Problems encountered

In the course of conducting the tests reported in this document, several problems with the LAP prototype were encountered. These problems pertain specifically to defects in the existing prototype, and would not necessarily recur if another prototype was built according to the same design. Nonetheless, we document these problems for the sake of completeness.

Sticky response. When the LAP identifies a direction of origin for a signal, it illuminates an LED light corresponding to that direction. Normally, the light remains illuminated for a short period (2 s) and then turns off. It sometimes happened, however, that an LED light would turn on and fail to turn off. In particular, LED 10 (when the LEDs are labelled 0 to 15) frequently exhibited this behaviour. In some cases, this problem would occur even when no stimulus had been presented to the array. When in this state, the array would continue to respond to stimuli presented from other directions. It could only be reset by powering down. Given the behaviour pattern of this problem, it seems likely that it originates in the user interface itself, perhaps as a result of a loose wire or poor connection.

Unresponsiveness. It also happened that sometimes the LAP would simply cease to respond to stimuli at all. In such cases, it was necessary to power down the array and restart it. This occurred at irregular intervals, but usually occurred at least once in a two–hour period. On one occasion it occurred so often that it prevented reliable data collection entirely. The origin of this problem is not known.

Dead channel. As was mentioned in Section **3.2.2** and illustrated in **Figure 11**, one of the output channels (Sensor 11) never responded in any of the tests. This channel must be considered dead. Whether this resulted from a problem with the sensor, the wiring, or the output LED for that channel is not known. The total number of responses elicited from the array during the resolution tests of that channel, all of which appeared in neighbouring channels, suggest that the problem is with the input signal, not the output. No obvious wiring problems were evident.

5 Comparison to related systems

In this Section, we briefly review the technical profile and performance characteristics of four related soldier–worn acoustic gunfire localization systems that are either commercially available or have been described in the literature. Comparing these systems to the LAP generates several recommendations for future work on this technology which are presented in Section **6**.

As was stated briefly in Section **1**, at least two commercial systems using shoulder–mounted acoustic sensors are available. The Shoulder-Worn Acoustic Targeting System (SWATS) [23] from QinetiQ North America has been selected by the United States Army for deployment in theatre [24]. The unit is small and light (0.45 kg), has low power requirements (< 1 W) for long operating periods, and low latency (< 1 s) for shot detection [23]. It reports both bearing and range to the shooter. Detailed performance characteristics, however, such as bearing and range accuracy, false positive rate, and suitability for urban environments do not appear to be publicly available.

Somewhat more information is available for the Boomerang Warrior–X system [25] from BBN Technologies. Also shoulder–mounted, this unit is light (0.34 kg), can operate for up to 12 hours without battery replacement, and has a fast (< 1 s) response time. Like SWATS, it provides the wearer with both bearing and range estimates to the shooter. The promotional materials for the device report a bearing accuracy of better than 7.5° and a range accuracy of ±20%, as well as false positive and false negative rates of < 2% and < 5%, respectively. The company also states that SWATS is effective under rapid fire, but no further details are provided. To our knowledge, these results have not been verified by a third party, but if true they indicate that the system provides more and better information than does the LAP in its present form.

An earlier wearable system consisting of helmet–mounted acoustic sensors was also developed by BBN Technologies [18]. Twelve microphones were mounted flush against each helmet, and the acoustic signals were combined with an orientation sensor and Global Positioning System (GPS) receiver to provide position data. The power consumption of each system was estimated at 25 W, which is too high for a deployed system but acceptable for a prototype. For a single system, they found that at least one third of shots could be localized in both azimuth and elevation to better than 5° and approximately three–quarters to better than 20°. The range to the shooter was estimated with ±20% precision in over 80% of cases. They detected over 90% of shots and were also able to identify the calibre of the bullet in over 80% of cases.

Each unit was also integrated with a radio transmitter and receiver, allowing the information gathered by each unit to be shared with other units. When data were shared and combined, the accuracy of the localization improved. Six helmet systems were distributed over an area approximately 100 m × 100 m in size, and the shooter, located approximately 200 m away, fired into the area occupied by the sensor network. Under these conditions, azimuth and elevation were determined to within 5° in 90% of cases. Range was determined with ±5% accuracy in 50% of cases, and all range estimates were accurate to within ±20%. Moreover, all shots were detected by the network, and the calibre of the bullet was correctly identified in all cases. These tests were carried out in open terrain with stationary sensors, so they did not test the robustness of the system under reverberant urban conditions nor under sensor motion.

Finally, we examine another wearable, networked system of acoustic sensors for shooter localization developed at Vanderbilt University (Nashville, Tennessee) [26]. In this case, each sensor node consisted of four acoustic sensors attached with fixed separation to a small computing mote. Each mote was also equipped with a three–axis digital compass to provide orientation, and was attached to a helmet. The mote communicated via a standard Bluetooth connection with the soldier's Personal Digital Assistant (PDA), which was assumed to have GPS capability for tracking position.

The motes were connected through a wireless network and shared information using a multi–hop routing protocol. All nodes in the network reported to a base node, which then shared the results of the localization algorithm back into the network. The system provided estimates for bearing, range, trajectory, calibre, and weapon type. Perhaps the most innovative feature of this system is its robustness under urban conditions: it does not require line–of–sight to the shooter and can tolerate multipath effects. Tested in an urban environment but without moving sensors, a system of ten sensors covering an area of 30 m × 30 m provided trajectory precision of 1° and identified both calibre and weapon (from a pool of six types) with over 95% accuracy. Range estimates were accurate to within 5% for close–range (50 m) shots but rose to 25% for far–range (300 m) shots; it is believed that the reasons for this degraded performance are understood and that range estimation can be improved. When the network was removed and single sensors operated independently, they still achieved bearing errors of roughly 1° and range errors of $5 - 10\%$. Shot detection rates averaged only about 40% for individual sensors; the network of sensors, by contrast, detected the shots in 96% of cases, which demonstrated the value of sensor data fusion in this context.

Comparison of these systems to the LAP highlights numerous deficiencies of the latter. Whereas the LAP provides only azimuthal bearing, other systems also provide some combination of range, elevation, trajectory, miss distance, shot calibre, and weapon type. These additional outputs are computable because these systems detect not only the firearm's muzzle blast but also the shockwave produced by the supersonic bullet. Because we have not yet carried out live–fire tests of the LAP, it is not known how it will respond to an incident shockwave, but certainly the system does not presently distinguish between shockwaves and muzzle blasts and so has no capacity to make productive use of the difference.

With respect to the precision of the azimuthal bearing estimates, the LAP is considerably less precise (see **Table 3**) than the systems described above. Indeed, the best case angular precision of the LAP is worse than the precision reported for other systems. The principal reason for this limitation is the direction–finding algorithm used by the current prototype: the direction to the shooter is indicated simply by indicating the sensor that best approximates the direction to the source of the muzzle blast's source. As such, the angular resolution is directly limited by the number of sensors in the array. A more precise bearing estimate could be achieved by using a Time Difference of Arrival (TDoA) algorithm that takes into account the relative time of arrival of the acoustic wave at neighbouring sensors. This should be considered a high priority in any future development of this technology; more recommendations for future work are given in the next Section.

6 Recommendations for future work

In this Section, we make a number of recommendations to guide any future development of this technology. Several of these recommendations echo those made by the contractors who developed the current prototype of the LAP [27], and several are based on the findings reported in this document.

Given the large number of recommended improvements and alterations that follow, and given the performance gap between the current LAP prototype and commercially available systems, it will be important to carefully consider whether or not further in–house research and development of this technology is warranted. The finest research–grade systems in this area [18, 26] are more advanced than commercial systems and include capabilities (such as effective localization in urban settings [26]) largely or wholly missing from currently available commercial systems. If the capabilities of commercial systems are deemed adequate for the current and forecasted needs of the Canadian Forces, there can be little reason to pursue a research program on this topic. On the other hand, if a more advanced system is desired, research on the topic is warranted provided it proceeds from an understanding of the leading achievements on this topic and produces a competetive system. It would be sensible for any further research and development of this technology within DRDC to draw on the expertise of DRDC Valcartier [20].

With that in mind, we present the following recommendations for future work:

Add shockwave detection for range, trajectory, and calibre estimation. It is desirable that any future system detect and distinguish between the conical shockwave produced by the supersonic projectile and the spherical acoustic blast wave produced at the muzzle of the firearm. Shockwaves are, in fact, more reliable than muzzle blasts because they are not damped by the shooter's distance (or use of a silencer) nor can they be spoofed (the pressure profile of a supersonic shockwave is very distinctive [30]). On the basis of shockwave detections, one can estimate the bullet trajectory, speed, and calibre [26]. If the muzzle blast is also detected, then the range to the shooter can be estimated based on the time difference of arrival between the shockwave and the muzzle blast [26].

Improve shockwave and muzzle blast detection and discrimination. In a system that bases its localization estimates on both shockwaves and muzzle blasts detections, it is essential that the two types of signal be cleanly distinguished from one another. Several methods for doing so have been proposed, including state machine models [26] and wavelet transforms [31, 32]. A robust detection method will likely also help to reduce both false positives and false negatives.

Improve angular resolution. The current system identifies the single sensor sector which best approximates the direction to the source of the impulsive stimulus. As such, the array's angular resolution is directly limited by the number of sensors. (For $n = 16$ sensors, the best case angular resolution is $360°/16 = 22.5°$; however, **Table 3** shows that in practice the resolution is somewhat worse.) It is imperative, however, that the resolution be improved by taking into account time–of–arrival differences between adjacent sensors. An angular precision of at least 5° should be targeted in order to be competitive with existing systems. Clearly, this would also involve replacing the current LED indicator lights with an interface capable of showing finer–grained directions.

Improve localization in urban settings. Urban environments pose special challenges for acoustic gunfire localization systems. Because of acoustic reflections from buildings, a single shot can appear to multiply itself. Such multipath effects can easily disrupt a system that localizes based on an acoustic wave's time of arrival. In an urban setting, the sensor may not even have a line–of–sight to the shooter. To our knowledge, only one group has described a method for overcoming these challenges [26, 33, 34, 35]. The method involves defining a consistency function over a set of sensor signals that quantifies the degree to which the signal set is consistent with a shot originating from a given space–time co–ordinate, and then performing an intelligent bisection search through the local space–time to find the point of greatest consistency. This group reports that it is able to identify the shooter's location to within less than 2 m, even in urban environments.

Reduce power consumption. A crucial requirement for any deployable direction–finding system is that its power requirements be low enough to allow long–term use with a portable power supply. The current prototype requires approximately 10 W of power, which is too high (see discussion of power supply below). There are several possible ways to reduce the power consumption:

- *Channel multiplexing.* In the current design, each of the sixteen sensors is sampled by a dedicated ADC. Power consumption could be reduced by reducing the number of ADCs through time–division multiplexing. In this scenario, a single ADC would sample a set of channels sequentially. Given an inter–channel switching time of t_s and a number of channels per ADC of n_c, the maximal possible sampling rate per channel would be $f_s^{max} = 1/(t_s n_c)$. For a reasonable value of the switching time ($t_s \simeq 200$ ns [27]), and supposing we wanted to multiplex all sixteen channels into one ADC ($n_c = 16$), then $f_s^{max} \simeq 300$ kHz. Given that the current sampling rate is $f_s \simeq 100$ kHz, this approach would be feasible.

  Multiplexing would also allow the number of amplifiers to be reduced. At present there is one amplifier per input channel, but in a multiplexing scenario amplification could be applied only to the multiplexed signal, resulting in considerable power savings.

A complication inherent to multiplexing is that instead of being sampled simultaneously, as in the current design, the channels would be sampled sequentially. This would introduce a systematic bias into the direction–finding. Note, however, that if the inter–channel sampling interval is $t_s \simeq 200$ ns, as suggested above, and if the speed of sound is roughly $v \simeq 340$ m/s [36], then the distance travelled by the wave–front while the sampler is switching channels would be just $d = vt_s \simeq 0.07$ mm, which is small compared to the inter–sensor spacing (~ 3 cm). As such, the bias resulting from the sequential sampling would probably be negligible. Even so, it should be possible to correct for the bias if it were thought necessary.

- *Sensor reduction.* A straight–forward method to reduce power consumption would be to reduce the number of sensors in the array. If that were the only change, it would result in a decrease in angular resolution, since each sensor's angular sector would be enlarged. It should be possible, according to results reported by others [18, 26], to restore and even improve the resolution by the TDoA method described above.

- *Stand–by mode.* Many electronic devices conserve power by having a low power stand–by mode in which functionality is reduced. The LAP, however, must constantly sample and analyze the auditory environment, which involves most of its components. It may be possible to revise the design so that only a portion of the system is continuously operational, with the detection of a shot serving as the signal to wake the entire device [26].

Add portable power supply. The current prototype is powered by a 10 V DC power supply. For the system to be of practical use, however, it would be necessary to power it from a light–weight portable battery.

Batteries are often rated according to their specific energy, which states the energy per unit mass that the battery can provide. Commercially available batteries have specific energies ranging from 40 $\text{W} \cdot \text{h/kg}$ (NiCd battery [37]) to 200 $\text{W} \cdot \text{h/kg}$ (Li–ion battery [38]). Rare battery types can have specific energies of up to 330 $\text{W} \cdot \text{h/kg}$ (LiS battery [39]).

Given a battery–type with a specific energy e and a battery pack with a mass m, a device that uses power P would be able to operate for a time given by $t = em/P$. If we wanted, therefore, a system that could operate for 10 hours with a 300 g Li–ion battery pack, we would require that the power requirements of the device be not more than $P = em/t = (200\ \mathrm{W}\cdot\ \mathrm{h/kg})(0.3\ \mathrm{kg})/(10\ \mathrm{h}) \simeq 7\ \mathrm{W}$. This is comparable to the current power consumption of 10 W. If some of the power reduction strategies outlined above were implemented, a battery–powered system would be feasible.

Improve sensitivity. As shown in **Figures 5** and **7**, the background noise level strongly affects the ability of the LAP to reliably identify the onset and direction of origin of an impulsive stimulus.

Problems were encountered in both low–noise and high–noise environments. In low noise, the LAP appeared to be very sensitive, often responding on multiple channels to a single stimulus. It is possible that the multiple responses were due to acoustic reflections from the walls of the laboratory. It would be desirable to study the input signals and processing steps in such cases to ascertain the cause. It is possible that a modification could be made to the signal processing in order to mitigate the problem.

In high noise, on the other hand, the LAP ceased to respond to stimuli. Although the precise noise level at which responses ceased depended on the stimulus being presented, it was found that the unaided ear could still reliably detect the stimulus in noise levels for which the LAP was unable to do so. Improving the ability of the LAP to detect stimuli in high–noise environments should be a priority. It may be possible to improve the performance at high–noise levels by adjusting parameters in the signal processing algorithms.

Improve reliability. As was discussed in Section **4**, several technical problems with the LAP recurred during the course of experimentation. The problem of unresponsiveness, in which the system stopped responding to stimuli, could be remedied by adding a "watchdog" circuit which would automatically reset the device if it failed to respond to a signal. A status indicator light would also be an advantageous addition.

Improve user interface. The current user interface, consisting of an array of LED lights, is inappropriate for a field–deployable system. Consideration must be given to the human factors problem of designing an intuitive and informative interface.

Commercial gunfire localization systems use a combination of visual and auditory cues to indicate the direction to the shooter [23, 25], the visual data being sent to a small dedicated display that clips to the soldier's gear and the auditory data to an earpiece. A circular visual interface, resembling a compass or a clock, would be an intuitive way to present the directional information. One could envisage a helmet with a Heads–Up Display onto which the directional information would be projected. Auditory signals would preferably be integrated with a multi–purpose communication earpiece.

Improve sensor mounting. As was illustrated in **Figure 12**, the mounting of the sensor array onto the helmet in the current prototype is rudimentary, and improving the mounting must be a central challenge for future designs.

For a helmet–mounted sensor array, one could imagine the sensors being either attached to the exterior of the helmet or embedded into the helmet itself. In the former case, a mounting technique that keeps the sensors flush against the surface of the helmet would be required. It has been suggested [27] that metal or plastic clips could be used for this purpose. One would face the problem, however, of how to route the wires (presumably also on the exterior of the helmet) to the processing unit. Such a system would be exposed to damage from blunt force and from the elements, and it would probably be impractical to make it robust. Another solution reported in the literature [18] could be to embed the sensors into a plastic overlay that fits over a helmet.

A design in which sensors were embedded into the helmet would be more robust. This possibility, however, would necessitate re–designing and re–fabricating the entire helmet, and it would be harder to repair, replace, and upgrade the localization system.

Although several other research groups have experimented with helmet–mounted sensors [18, 26], commercial vendors have opted instead for small shoulder–mounted units [23, 25].

Distribute array. A more radical modification of the LAP would be to distribute it spatially over a larger area [18, 26]. In particular, the number of sensors worn by an individual soldier could be reduced if the sensors could share their data with one another over a wireless network. In this way, the power requirements for each soldier would be reduced and, because the sensors would cover a larger spatial region, more precise identification of the source direction would be possible. Major modifications to the direction–finding algorithms would be necessary since the sensor array would have a flexible, not fixed, geometry.

This model would require that certain technological capabilities be present to support the localization process. The position of each sensor would need to be known, likely requiring that the sensors be integrated with GPS receivers (although other solutions may also be possible [40]). In addition, all signal inputs to the direction–finding algorithm would need to be synchronized in time, which is a challenging requirement [41, 42, 43, 44]. Finally, it would be necessary for the distributed sensors to communicate with one another. This could be accomplished over a Mobile Ad Hoc Network (MANET) [45, 46], if one was present, using a standard or customized [47] protocol.

Given the advantages to be gained from a distributed wireless sensor network, this approach to the problem of gunfire localization may warrant further study.

7 Conclusions

In this document, we have described several experiments conducted to evaluate the performance characteristics of a prototype LAP, a helmet–mounted acoustic sensor array for identifying the direction of origin of short duration impulsive sounds such as gunfire or explosions.

The LAP consists of a sensor array, a signal processing unit, a user interface, and a power supply. The sensor array is a roughly circular set of acoustic sensors affixed around the perimeter of an army helmet. The signal processing unit samples the sensor outputs, monitoring the acoustic environment for the advent of a short duration impulse such as is associated with a small–arms muzzle blast or explosive. When detected, the system identifies the direction from which the sound originated, and outputs the result to a simple display consisting of an array of LED lights. The current prototype, lacking a battery pack, is not portable.

Several experiments were carried out on the LAP. Its response in backgrounds of stationary noise was studied. As expected, its performance degraded as noise levels increased, but it was also found that it degraded faster than expected. Tests of the angular resolution of the direction–finding algorithm showed that it correctly identified the direction of origin in approximately 50% of cases. Because of the direction–finding algorithm that is implemented, however, the angular resolution is poor compared to similar systems described in the literature. Finally, the recovery time of the LAP in a rapid–fire scenario was studied, and it was found that the system can detect a stimulus and refresh to a listening state in a few hundred milliseconds.

A number of recommendations have been made for further development of this technology. The chief objectives of any further work should be to provide more information, such as range to the shooter and bullet calibre, improve the angular resolution, improve performance under multipath effects such as occur in urban settings, reduce power consumption, provide a portable power supply, improve the system's performance in high–noise environments, improve and secure the mounting of the sensors, and modify the user interface. Finally, a proposal whereby the idea of the LAP could be applied to a flexible sensor array distributed over a larger spatial area has been presented as a possible alternate route for future developments.

This page intentionally left blank.

Annex A: Generating multi–channel audio files

All of the experiments described in this document were carried out by playing multi–channel audio files through a set of loudspeakers. These multi–channel audio files were generated prior to each experiment. In this Annex, we give a brief description of the technical process of generating such files.

The advantage of multi–channel audio files is that they allow up to eight independent audio signals to be combined into one file and played using an ordinary media player, such as (in our case) Windows Media Player. This is an attractive alternative to using an expensive sound–mixing software suite.

There are two steps to generate a multi–channel audio file. First, the component single–channel (mono) files must be assembled. These files can be of any length, but they must all contain the same number of samples. They must also have the same sampling rate to avoid distortion. In our experiments, these mono files were generated in MATLAB, and were in the `.wav` format.

To combine the mono files into a single multi–channel file, we use a Microsoft command–line utility called `wavavimux` [48]. This utility converts a set of mono `.wav` files into a multi–channel audio file in the `.avi` format. The command to be used is

```
wavavimux -o [output file] -iwav 8 [input files] -mask 255    (A.1)
```

where the output and input filenames must include both the paths and the extensions. The `mask` value specifies the channels to which the input files should be written. In this example, we assume that eight input files are used, so that all of the channels are active (255 = 11111111 in binary notation). If fewer input files were used, zeros in the binary mask value would indicate which channels in the output file should be blanked out.

The `wavavimux` utility outputs a file that satisfies the Surround Sound 7.1 standard [49], which means that a gain is automatically applied to each of the channels in the multi–channel file. To compensate for this fact, the gain on each of the channels was individually modified in the multi–channel amplifier to ensure a uniform sound level from each of the loudspeakers. The Surround Sound 7.1 standard also reserves Channel 4 for a sub–woofer loudspeaker and applies a low–pass filter to the audio in that channel. Since we were not using a sub–woofer in our experiments, we consistently excluded Channel 4.

References

[1] Atherley, G. and Noble, W. (1970), Effect of ear-defenders (ear-muffs) on the localization of sound, *British Journal of Industrial Medicine*, 27(3), 260–265.

[2] Noble, W. and Russell, G. (1972), Theoretical and practical implications of the effects of hearing protection devices on localization ability, *Acta Oto-laryngologica*, 74(1), 29–36.

[3] Noble, W., Murray, N., and Waugh, R. (1990), The effects of various hearing protectors on sound localization in the horizontal and vertical planes, *American Industrial Hygiene Association Journal*, 51(7), 370–377.

[4] Abel, S. and Armstrong, N. (1993), Sound localization with hearing protectors, *Journal of Otolaryngology*, 22(5), 357–363.

[5] Vause, N. and Grantham, D. (1999), Effects of earplugs and protective headgear on auditory localization ability in the horizontal plane, *Human Factors*, 41(2), 282–294.

[6] Bolia, R., D'Angelo, W., Mishler, P., and Morris, L. (2001), Effects of hearing protectors on auditory localization in azimuth and elevation, *Human Factors*, 43(1), 122–128.

[7] Brungart, D., Kordik, A., Simpson, B., and McKinley, R. (2003), Auditory localization in the horizontal plane with single and double hearing protection, *Aviation, Space, and Environmental Medicine*, 74(9), 937–946.

[8] Abel, S. and Paik, J. (2005), Sound source identification with ANR earmuffs, *Noise & Health*, 7(27), 1–10.

[9] Lukas, K. and Ahroon, W. (2006), Free-field sound localization with nonlinear hearing protective devices, *Journal of the Acoustical Society of America*, 120(5), 3080–1.

[10] Abel, S., Tsang, S., and Boyne, S. (2007), Sound localization with communication headsets: Comparison of passive and active systems, *Noise & Health*, 9(37), 101–107.

[11] Carmichel, E., Harris, F., and Story, B. (2007), Effects of binaural electronic hearing protectors on localization and response time to sounds in the horizontal plane, *Noise & Health*, 9(37), 83–95.

[12] Borg, E., Bergkvist, C., and Bagger-Sjöbäck, D. (2008), Effect on Directional Hearing in Hunters Using Amplifying (Level Dependent) Hearing Protectors, *Otology & Neurotology*, 29(5), 579–585.

[13] Torrey Pines Logic, Mirage 1200 (online), `http://www.tplogic.com/products/mirage1200.html` (Access Date: 13 April 2011).

[14] Rheinmetall Defence, Sniper Locating System (online), `http://www.rheinmetall-detec.de/index.php?fid=4235&lang=3&pdb=1` (Access Date: 13 April 2011).

[15] Radiance Technologies, WeaponWatch (online), `http://www.radiancetech.com/products/weaponwatch.htm` (Access Date: 13 April 2011).

[16] M2 Technologies, Anti-Sniper Infrared Targeting System (online), `http://www.m2tech.us/company/accomplishments.html` (Access Date: 13 April 2011).

[17] Duckworth, G., Gilbert, D., and Barger, J. (1997), Acoustic counter-sniper system, In *Proceedings of SPIE International Symposium on Enabling Technologies for Law Enforcement and Security*, pp. 262–275.

[18] Duckworth, G., Barger, J., Carlson, S., Gilbert, D., Knack, M., Korn, J., and Mullen, R. (1999), Fixed and wearable acoustic counter-sniper systems for law enforcement, In *Proceedings of SPIE International Symposium on Sensors, C3I, Information, and Training Technologies for Law Enforcement*, pp. 210–230.

[19] ShotSpotter, Solutions (online), `http://www.shotspotter.com/solutions` (Access Date: 13 April 2011).

[20] Bédard, J. and Paré, S. (2003), Ferret, a small-arms' fire detection system: Localization concepts, In *Proceedings of SPIE International Symposium on Sensors and C3I Technologies for Homeland Defense and Law Enforcement II*, pp. 497–509.

[21] US Army, C., Project Manager, Gunfire Detection System (GDS) (online), `http://www.pica.army.mil/pmccs/d5forceapplication/p5_2special/gds.htm` (Access Date: 14 April 2011).

[22] AAI Textron Systems, Projectile Detection and Cueing (PDCue) (online), `http://www.aaicorp.com/pdfs/aai_pdcue02-14-08b.pdf` (Access Date: 14 April 2011).

[23] QinetiQ North America, Ears SWATS VMS FSS Gunshot Localization System (online), `http://www.qinetiq-na.com/products-survivability-ears.htm` (Access Date: 14 April 2011).

[24] QinetiQ North America, QinetiQ's Gunfire Detection System is U.S. Military Solution of Choice (online), `http://www.prweb.com/releases/gunfire-detection/swats/prweb8050380.htm` (Access Date: 14 April 2011).

[25] Raytheon BBN Technologies, Boomerang Warrior-X: Compact Soldier Worn Shooter Detection System (online), `http://bbn.com/resources/pdf/Boomerang-Warrior-X-072210.pdf` (Access Date: 14 April 2011).

[26] Völgyesi, P., Balogh, G., Nádas, A., Nash, C., and Lédeczi, A. (2007), Shooter localization and weapon classification with soldier-wearable networked sensors, In *Proceedings of the 5th international conference on mobile systems, applications and services*, MobiSys '07, pp. 113–126, New York, NY, USA: ACM.

[27] Haykin, S., Wiklund, K., Freibert, A., and Zhang, F. (2010), Improving Situational Awareness in Noisy Environments: A helmet-based system for speech enhancement, hearing protection, and shock localization, (DRDC Toronto CR 2010-048) Defence R&D Canada – Toronto.

[28] M-Audio, Delta 1010LT: 10-in/10-out PCI Virtual Studio (online), `http://www.m-audio.com/products/en_us/Delta1010LT.html` (Access Date: 23 November 2010).

[29] BiAmp Professional Audio Systems, MCA Series Amplifier (online), `http://www.biamp.com/products/amplifiers/mca.aspx` (Access Date: 23 November 2010).

[30] Stoughton, R. (1997), Measurements of small-caliber ballistic shock waves in air, *Journal of the Acoustical Society of America*, 102(2), 781–787.

[31] Sadler, B., Pham, T., and Sadler, L. (1998), Optimal and wavelet-based shock wave detection and estimation, *Journal of the Acoustical Society of America*, 104, 955–963.

[32] Mays, B. (2001), Shockwave and muzzle blast classification via joint time frequency and wavelet analysis, (Technical Report 20783-1197) Army Research Lab, Alelphi, MD.

[33] Simon, G., Maróti, M., Lédeczi, A., Balogh, G., Kusy, B., Nádas, A., Pap, G., Sallai, J., and Frampton, K. (2004), Sensor network-based countersniper system, In *Proceedings of the 2nd international conference on embedded networked sensor systems*, SenSys '04, pp. 1–12, New York, NY, USA: ACM.

[34] Lédeczi, A., Nádas, A., Völgyesi, P., Balogh, G., Kusy, B., Sallai, J., Pap, G., Dóra, S., Molnár, K., Maróti, M., and Simon, G. (2005), Countersniper system for urban warfare, *ACM Transactions on Sensor Networks*, 1, 153–177.

[35] Lédeczi, A., Völgyesi, P., Maróti, M., Simon, G., Balogh, G., Nádas, A., Kusy, B., Dóra, S., and Pap, G. (2005), Multiple simultaneous acoustic source localization in urban terrain, In *Proceedings of the 4th international symposium on information processing in sensor networks*, IPSN '05, Piscataway, NJ, USA: IEEE Press.

[36] Everest, F. (2001), The Master Handbook of Acoustics, McGraw-Hill.

[37] Eveready Battery Company, Energizer No. CH35 Engineering Data (online).

[38] Panasonic Corporation, Lithium-Ion Batteries: Individual Data Sheet: CGR18650E (online), `http://www.panasonic.com/industrial/includes/pdf/Panasonic_LiIon_CGR18650E.pdf` (Access Date: 23 November 2010).

[39] Sion Power, Lithium Sulfur Rechargeable Battery Data Sheet (online), `http://sionpower.com/pdf/articles/LIS%20Spec%20Sheet%2010-3-08.pdf` (Access Date: 23 November 2010).

[40] Amundson, I. and Koutsoukos, X. (2009), A survey on localization for mobile wireless sensor networks, In *Proceedings of the 2nd international conference on mobile entity localization and tracking in GPS-less environments*, MELT'09, pp. 235–254, Berlin, Heidelberg: Springer-Verlag.

[41] Maróti, M., Kusy, B., Simon, G., and Lédeczi, A. (2004), The flooding time synchronization protocol, In *Proceedings of the 2nd international conference on embedded networked sensor systems*, SenSys '04, pp. 39–49, New York, NY, USA: ACM.

[42] Younis, O. and Fahmy, S. (2005), A scalable framework for distributed time synchronization in multi-hop sensor networks, In *Second Annual IEEE Communications Society Conference on Sensor and Ad Hoc Communications and Networks*, pp. 13 – 23.

[43] Sallai, J., Kusy, B., Lédeczi, A., and Dutta, P., On the Scalability of Routing Integrated Time Synchronization, In *Proceedings of the European Workshop Wireless Sensor Networks (EWSN '06)*.

[44] Damarla, T., Kaplan, L., and Whipps, G. (2010), Sniper Localization Using Acoustic Asynchronous Sensors, *Sensors Journal, IEEE*, 10(9), 1469 –1478.

[45] Basagni, S., Conti, M., Giordano, S., and Stojmenovic, I. (2004), Mobile Ad Hoc Networking, Wiley-IEEE Press.

[46] Genik, L., Salmanian, M., Mason, P., Schotanus, H. A., Verkoelen, A., and Hansson, E. (2004), Mobile Ad Hoc Security from a Military Perspective, (DRDC Ottawa TR 2004-252) Defence R&D Canada – Toronto.

[47] Maróti, M. (2004), Directed flood-routing framework for wireless sensor networks, In *Proceedings of the 5th ACM/IFIP/USENIX International Conference on Middleware*, pp. 99–114.

[48] Microsoft Corporation, Creating 7.1 Audio (online), `http://www.microsoft.com/windows/windowsmedia/howto/articles/creating71audio.aspx` (Access Date: 1 November 2010).

[49] Crutchfield (Tara Wisnewski), Understanding Surround Sound Formats (online), `http://www.crutchfield.com/S-uaCT5J7oKCd/learn/learningcenter/home/hometheater_surround.html` (Access Date: 27 July 2011).

Acronyms

ADC	Analog-to-Digital Converter
CPLD	Complex Programmable Logic Device
DARPA	Defense Advanced Research Projects Agency
dB	Decibels
dBA	A-weighted Decibels
DRDC	Defence Research and Development Canada
DSP	Digital Signal Processor
GPS	Global Positioning System
LAP	Localization Array Processor
LED	Light Emitting Diode
PDA	Personal Digital Assistant
SNR	Signal to Noise Ratio
SPL	Sound Pressure Level
SWATS	Shoulder-Worn Acoustic Targeting System
TDoA	Time Difference of Arrival

This page intentionally left blank.

Distribution list

DRDC Toronto TM 2011-062

Internal distribution

2	Author
1	Stergios Stergiopoulos, DRDC Toronto
1	Library, DRDC Toronto

Total internal copies: 4

External distribution

1	Library and Archives Canada

Total external copies: 1

Total copies: 5

This page intentionally left blank.

DOCUMENT CONTROL DATA

(Security classification of title, body of abstract and indexing annotation must be entered when document is classified)

1. ORIGINATOR (The name and address of the organization preparing the document. Organizations for whom the document was prepared, e.g. Centre sponsoring a contractor's report, or tasking agency, are entered in section 8.)

Defence R&D Canada – Toronto
1133 Sheppard Avenue West, PO Box 2100, Toronto
ON M3M 3B9, Canada

2. SECURITY CLASSIFICATION (Overall security classification of the document including special warning terms if applicable.)

UNCLASSIFIED
(NON-CONTROLLED GOODS)
DMC A
REVIEW: GCEC JUNE 2010

3. TITLE (The complete document title as indicated on the title page. Its classification should be indicated by the appropriate abbreviation (S, C or U) in parentheses after the title.)

Evaluation of a Helmet–Mounted Auditory System for Gunfire Localization

4. AUTHORS (Last name, followed by initials – ranks, titles, etc. not to be used.)

Burrell, C.; Saunders, D.

5. DATE OF PUBLICATION (Month and year of publication of document.)

December 2011

6a. NO. OF PAGES (Total containing information. Include Annexes, Appendices, etc.)

56

6b. NO. OF REFS (Total cited in document.)

49

7. DESCRIPTIVE NOTES (The category of the document, e.g. technical report, technical note or memorandum. If appropriate, enter the type of report, e.g. interim, progress, summary, annual or final. Give the inclusive dates when a specific reporting period is covered.)

Technical Memorandum

8. SPONSORING ACTIVITY (The name of the department project office or laboratory sponsoring the research and development – include address.)

Defence R&D Canada – Toronto
1133 Sheppard Avenue West, PO Box 2100, Toronto ON M3M 3B9, Canada

9a. PROJECT OR GRANT NO. (If appropriate, the applicable research and development project or grant number under which the document was written. Please specify whether project or grant.)

14dc02

9b. CONTRACT NO. (If appropriate, the applicable number under which the document was written.)

10a. ORIGINATOR'S DOCUMENT NUMBER (The official document number by which the document is identified by the originating activity. This number must be unique to this document.)

DRDC Toronto TM 2011-062

10b. OTHER DOCUMENT NO(s). (Any other numbers which may be assigned this document either by the originator or by the sponsor.)

11. DOCUMENT AVAILABILITY (Any limitations on further dissemination of the document, other than those imposed by security classification.)

(X) Unlimited distribution
( ) Defence departments and defence contractors; further distribution only as approved
( ) Defence departments and Canadian defence contractors; further distribution only as approved
( ) Government departments and agencies; further distribution only as approved
( ) Defence departments; further distribution only as approved
( ) Other (please specify):

12. DOCUMENT ANNOUNCEMENT (Any limitation to the bibliographic announcement of this document. This will normally correspond to the Document Availability (11). However, where further distribution (beyond the audience specified in (11)) is possible, a wider announcement audience may be selected.)

UNLIMITED

13. ABSTRACT (A brief and factual summary of the document. It may also appear elsewhere in the body of the document itself. It is highly desirable that the abstract of classified documents be unclassified. Each paragraph of the abstract shall begin with an indication of the security classification of the information in the paragraph (unless the document itself is unclassified) represented as (S), (C), or (U). It is not necessary to include here abstracts in both official languages unless the text is bilingual.)

The Localization Array Processor (LAP) is a system for acoustic localization of sources of rapid onset, short duration impulsive sounds, such as result from gunfire or explosions. The auditory environment is monitored by an array of helmet–mounted microphones. We report the results of a series of laboratory experiments to assess the performance characteristics of a prototype version of the LAP.

Three tests of the system were carried out. First, its operation in backgrounds of stationary noise was studied. It was found that the system's performance degraded rapidly as background noise levels increased, such that moderate levels of noise prevented the system from consistently identifying impulsive sounds in the environment. Second, tests of the direction–finding capability showed that the system was able to correctly identify the direction to the sound source in nearly 50% of cases, but that the angular resolution was not competitive with similar systems described in the literature. Finally, the recovery time of the system in a rapid fire scenario was studied, and it was found that it could detect an impulsive sound and refresh to a listening state in a few hundred milliseconds.

In general, the system's performance was satisfactory. Under optimal conditions it succeeded in identifying the azimuthal direction of origin of impulsive sounds, though the angular resolution provided by the current direction–finding algorithm is too coarse. The performance degraded in the presence of background noise. Numerous recommendations are made to guide future work in this area.

14. KEYWORDS, DESCRIPTORS or IDENTIFIERS (Technically meaningful terms or short phrases that characterize a document and could be helpful in cataloguing the document. They should be selected so that no security classification is required. Identifiers, such as equipment model designation, trade name, military project code name, geographic location may also be included. If possible keywords should be selected from a published thesaurus. e.g. Thesaurus of Engineering and Scientific Terms (TEST) and that thesaurus identified. If it is not possible to select indexing terms which are Unclassified, the classification of each should be indicated as with the title.)

Signal Processing; Situational Awareness; Acoustic Gunfire Localization; Sniper Detection

Defence R&D Canada

Canada's Leader in Defence
and National Security
Science and Technology

R & D pour la défense Canada

Chef de file au Canada en matière
de science et de technologie pour
la défense et la sécurité nationale

DEFENCE DÉFENSE

www.drdc-rddc.gc.ca