Proceedings of Australasian Conference on Robotics and Automation, 2-4 Dec 2014, The University of Melbourne, Melbourne, Australia

Audio Signalling as a Backup Communication Channel for Multi-Robot Systems

Gavin Suddrey, Pearl Gariano, Sam Cunningham-Nelson, Daniel Richards, Frederic Maire
School of Electrical Engineering and Computer Science
Queensland University of Technology, Australia
g.suddrey@qut.edu.au, f.maire@qut.edu.au

Abstract

This paper presents a low-bandwidth multi-robot communication system designed to serve as a backup communication channel in the event a robot suffers a network device fault. While much research has been performed in the area of distributing network communication across multiple robots within a system, individual robots are still susceptible to hardware failure. In the past, such robots would simply be removed from service, and their tasks re-allocated to other members. However, there are times when a faulty robot might be crucial to a mission, or be able to contribute in a less communication intensive area. By allowing robots to encode and decode messages into unique sequences of DTMF symbols, called words, our system is able to facilitate continued low-bandwidth communication between robots without access to network communication. Our results have shown that the system is capable of permitting robots to negotiate task initiation and termination, and is flexible enough to permit a pair of robots to perform a simple turn taking task.

1 Introduction

Multi-robot systems are of increasing interest to researchers in the field of robotics. This interest is due in large part to the range of advantages they provide over traditional single-robot systems.

A key advantage of multi-robot systems over their single-robot counterparts is an increase in system robustness [Parker and Fang, 2006]. Through the introduction of redundancy, multi-robot systems are designed to be robust against the loss of individual members of the system. This robustness, combined with task parallelisation, increased area coverage, and the ability to complete tasks too complex for a single robot [Su *et al.*, 2013], are what make multi-robot systems suitable to roles such as search and rescue [Couceiro *et al.*, 2013] and agriculture [Conesa-Munoz *et al.*, 2012].

Figure 1: A pair of humanoid NAO robots communicating by encoding messages into sequences of DTMF tones. These tones can then be transmitted between robots as audio.

According to [Basu and Redi, 2004], the success of robots in communicating and coordinating effectively is inherently linked to access to networking technologies. However, this reliance on networking technologies is also a key weakness in the application of multi-robot systems. Centralised systems, in which communications and resource management are managed by a single machine, are a prime example of where a single network failure can potentially neutralise the entire system [Couceiro *et al.*, 2013]. To address this issue, systems that are designed for communication robustness avoid the use of centralisation, instead opting to distribute communication and resource management across the entire system [Basu and Redi, 2004; Fang and Parker, 2005; Rohrmüller *et al.*, 2012; Couceiro *et al.*, 2013]. However, while these distributed systems are safe against single point of failure problems, individual robots can still be isolated from the system as a result of a network device failure.

To address the issue of a network device failure removing individual robots from the system, this paper presents a system that through audio signalling, is capable of permitting continued low-bandwidth communication. The use of audio as our communication medium was inspired by both the low cost, and ready availability of microphone and speaker technology. The contributions of this paper are as follows:

- First, we present a novel audio-based communication system, that by encoding messages into unique sequences of DTMF symbols called words, is capable of facilitating low-bandwidth communication between pairs of robots.
- Second, we demonstrate the reliability of the described system in transmitting low-bandwidth messages between robots within an office environment.
- Lastly, we demonstrate the ability of the system to facilitate basic multi-robot collaboration between homogeneous and heterogeneous pairs of robots.

The remainder of this paper is laid out as follows. Section 2 explores the background and related work. Section 3 describes the communication system. Section 4 describes the experimental setup and Section 5 catalogues the results. Finally, we present a discussion of our work in Section 6 and our conclusions in Section 7.

2 Background and Related Work

The issue of robustness has been of great interest to multi-robot researchers, and is a key area of investigation in much of the early work. ALLIANCE, presented in [Parker, 1998] is an early example of a multi-robot system designed with a focus on robustness. By using motivational concepts such as impatience and acquiescence, robots select the appropriate actions required for a task. Using these motivations, robustness is introduced by allowing robots to side-line faulty members if they are taking too long to complete a task.

However, while removing faulty members allows the system to redistribute tasks to competent members, there are times when a faulty member might still be capable of providing some overall benefit to the system. This might include instances where a faulty robot is the only one equipped with a particular tool, or where re-assignment to simplified tasks might be possible. RACHNA, a market-based multi-robot allocation framework presented in [Vig and Adams, 2006] addresses this issue by treating each capability of a robot as an independently auctionable service. A robot experiencing a fault in a sensor or actuator is thus able to simply remove the services affected by the fault from further auctions, while continuing to provide access to non-affected services.

Though multi-robot systems capable of handling a sensor or actuator fault without necessarily removing the affected robot from the system, these systems are still constrained by their reliance on network technology [Couceiro *et al.*, 2011]. This reliance on network communication means that any robot experiencing a network fault would, at present, simply vanish from the system, along with any useful capabilities it may provide. In order to address this issue, a backup method of communication, such as audio signalling, would be beneficial in providing a faulty robot with a means to either exit gracefully from the system, or continue to contribute in a less communication intensive role.

The use of audio signalling as a means of simple multi-robot communication was previously investigated in [Karimian *et al.*, 2006]. Through the use of audio-based DTMF signalling and sound localisation, robots in the system are tasked with mapping out and navigating to a set of shared targets. In the event that a robot finds a previously unknown target, it emits a recurring location-specific DTMF tone for a period of time. This tone communicates information on both the direction and name of the location which can then be utilized by any robot within hearing distance.

Extending on the concept of using audio communication to aide in localisation, [Eoh *et al.*, 2013] present a method for detecting and navigating to faulty robots. Through the use of a constantly emitted audio tone a faulty robot provides rescuer robots with a means of triangulating its location. These robots can then home in on and remove the faulty robot from the environment.

While audio transmissions comprised of a single tone might provide sufficient flexibility for communications needing only small dictionaries, it constrains the system to the available number of tones. To be able to facilitate more complex interactions between robots, a larger dictionary of terms is required. Through encoding messages into variable length sequences of tones, larger dictionaries become possible.

The use of sequences of DTMF tones to communicate complex information between robots has been explored in the work of [Schulz, 2009]. With the goal of creating a shared lexicon of place names, robots were encouraged to explore an environment and work collectively to build a shared vocabulary, using DTMF as the means of communication. This would involve the robots using audio communication to negotiate whether a location was novel, who would be responsible for naming it, as well as the name of the location. However, [Schulz *et al.*, 2011] reduce the use of DTMF communication to only querying for potential listeners. Once a listener had been identified, exchanging place names would then be handled by wireless communication.

3 The Communication System

While multi-robot systems are increasingly of interest to robotics researchers, the reliance of current systems on network communications represents a potential point of failure. The ability for robots to coordinate effectively without the need of networking communication however is non-trivial. The following sections describes an audio-based communication system that through the use of the DTMF symbol alphabet, permits pairs of robots to collaborate effectively without the need for network communication technology.

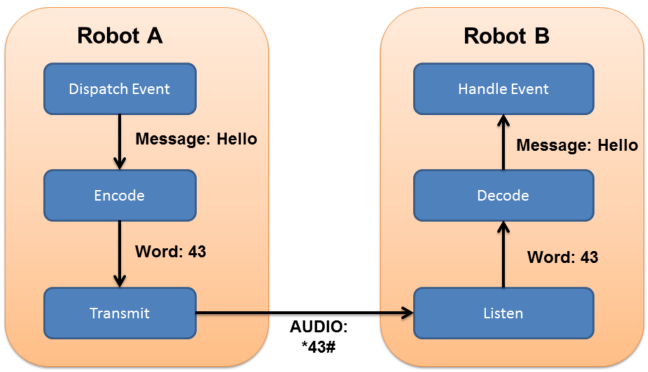

Figure 2: Messages can be passed between robots by first encoded them into words, comprised of DTMF symbols. Each symbol represents a tone that can be transmitted via audio.

3.1 DTMF Symbol Alphabet

Dual-Tone Multi-Frequency (DTMF) signalling is a well established signalling mechanism used in the field of telephony. Recognisable as the tones triggered by a telephone key press, each tone is formed from a mixture of two signals taken from mutually exclusive frequency groups, a high frequency signal group, and a low frequency signal group [International Telecommunication Union, 1990]. The resulting wave forms a distinct tone, representing a distinct symbol within the DTMF alphabet (see Figure 3).

	1209 Hz	1336 Hz	1477 Hz	1633 Hz
697 Hz	1	2	3	A
770 Hz	4	5	6	B
852 Hz	7	8	9	C
941 Hz	*	0	#	D

Table 1: Symbols within DTMF signalling alphabet.

The DTMF alphabet as defined in [International Telecommunication Union, 1990] provides for a total of 16 distinct symbols, with each symbol representing a transmittable audio tone (see Table 1). For the purpose of this system, 14 of these symbols are used in the construction of words, while two symbols are reserved as framing flags. These framing flags indicate the beginning and end of a transmission, and are known as the header and footer flags respectively.

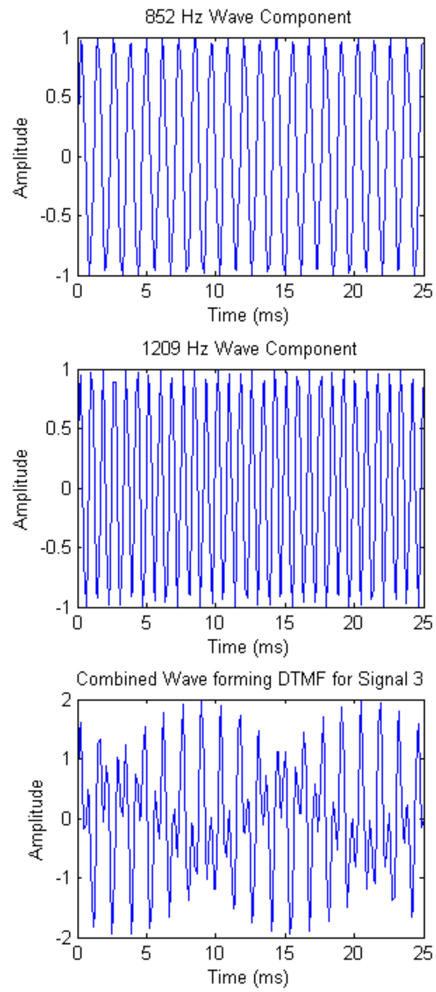

Figure 3: The DTMF tone representing the value of 3 is a mixture of an 852Hz and 1209Hz wave.

3.2 Word Structure

By combining a variable number of symbols from the previously described DTMF alphabet (excluding the framing flags), unique sequences of symbols called words can be created (see Table 2). These words can then be used to represent a wide variety of messages. Each symbol contained within a word can then be expressed as audio via its accompanying DTMF tone.

To map messages to words, a simple one-to-one mapping is applied. This mapping, if consistent throughout

Header	Word			Footer
*	4	2	3	#

Table 2: A word comprised of several DTMF symbols. The header and footer framing flags are applied during transmission of the word and are used to indicate the bounds of the transmission.

the system, allows each robot to understand the transmissions of any other robot within the system (see Figure 2). It is important to note that while this paper uses a one-to-one mapping, it is possible to create multiple words capable of expressing a single message (synonyms).

Message	Word
message_received	0
ready	1
finished	2
task_dock	41
task_follow	43

Table 3: An example dictionary. Robots that share the same dictionary are capable of communicating with each other.

As it is possible for a single instance of an emitted DTMF tone to be detected as multiple symbols representing that tone by the receiver, due to overlaps in tone emissions and detections, care must be taken to prevent ambiguity in the choice of words. For instance, the transmitted word "1" could be detected by the listener as either "1", "11" or even "111". Therefore, to avoid ambiguity, words can not be constructed using successive instances of the same symbol (e.g. "121" is valid while "112" is invalid). This allows the system to treat multiple successive detections of the same tone as a single emittance of that tone.

3.3 Word Transmission

This section describes a simple algorithm (see Algorithm 1) that given a message, selects, frames and transmits the corresponding word.

In order to transmit a message a corresponding word must first exist within the robot's dictionary. If a word exists, it is extracted from the dictionary and the robot prepares for transmission. However, in the event a word does not exist for the supplied message, an exception is raised to indicate failure. It is then up to the caller to determine how to handle this failure.

To prevent the robot from detecting and responding to its own transmissions, the algorithm will then block the ability of the speaking robot to detect.

Algorithm 1 Word Transmission Algorithm

```
function TRANSMIT(message)
    if message not index in words then
        raise exception
    end if
    word ← words(message)
    LISTENER.BLOCK
    PLAY(header_flag)
    for each tones in word do
        PLAY(tone)
    end for
    PLAY(footer_flag)
    LISTENER.RELEASE
end function
```

The robot will then initiate transmission by emitting the tone for the header flag symbol. This informs any listening robots that a word is about to be transmitted. The sequence of symbols comprising the word are then individually converted into their respective tones and emitted one at a time. The robot will then indicate that the transmission has ended by emitting the tone for the footer flag symbol.

Finally, the robot will re-enables its ability to detect words transmitted by other robots.

In the event that a transmission was successful, the listening robot will transmit an acknowledgement to signify that the transmission was received. The transmitting robot can then use this acknowledgement or the lack of an acknowledgement to determine its next course of action .

3.4 Word Detection

This section describes a set of algorithms (see Algorithms 2 and 3) for detecting transmit tones and for using the symbols representing these tones to reconstruct and decode transmitted words.

Algorithm 2 DTMF Tone Detection Algorithm

```
function DTMF_DETECT(wave, threshold)
    magnitudes ← GOERTZEL_MAGNITUDE(wave)
    high_group ← magnitudes(0 : 3)
    low_group ← magnitudes(4 : 7)
    high_group ← NORMALISE(high_group)
    low_group ← NORMALISE(low_group)
    high_index ← INDEX_OF_MAX(high_group)
    low_index ← INDEX_OF_MAX(low_group)
    if high_group(high_index) < threshold or
        low_group(low_index) < threshold then
        return null
    end if
    return GET_SYMBOL(high_index, low_index)
end function
```

In order to determine the likelihood that a DTMF frequency exists within a given audio sample, we first extract the Fourier coefficients for each of the frequencies used to construct DTMF tones. As the number of frequencies used to construct DTMF tones is small (see Table 1), the Goertzel algorithm [Goertzel, 1958] provides an efficient means for extracting the desired Fourier coefficients from the detected audio sample.

The coefficients extracted from the audio sample are then separated the high and low frequency groups. Each group is then normalised to obtain a ratio between the likelihood of a signal and background noise.

The index of the highest coefficient in each group is then extracted. The normalised coefficients at these indices are then compared against a threshold to determine if a valid tone has been detected. The indices of these coefficients can then be passed to a look-up table to extract the symbol representing that tone.

Using the symbol obtained in the previous algorithm, Algorithm 3 will then attempt to reconstruct the transmitted word, and decode it into its corresponding message.

Algorithm 3 Word Detection Algorithm

```
procedure LISTENER
    last_symbol ← NULL
    while true do
        wave ← AUDIO_CAPTURE
        symbol ← DTMF_DETECT(wave)
        if symbol ≠ null then
            if symbol = last_symbol then
                continue
            end if
            if symbol = header_flag then
                WORD_OPEN(word)
            else if symbol = footer_flag and
                    IS_OPEN(word) then
                if word in words then
                    WORD_CLOSE(word)
                    if WORD_VALID(word) then
                        msg ← DECODE_WORD(word)
                        PUBLISH(msg)
                    end if
                end if
            else if IS_OPEN(word) then
                WORD_APPEND(word, symbol)
            end if
            last_symbol ← symbol
        end if
        if IS_OPEN(word) and EXPIRED(word) then
            WORD_CLOSE(word)
        end if
    end while
end procedure
```

In the event a header flag symbol is obtained, a new word is opened for writing and its previous contents are erased. Subsequently obtained symbols are then appended to the message until a footer symbol is obtained.

On obtaining the footer symbol, the word is closed to prevent further writing, and the contents are checked for validity. The valid word can then be used to obtain its corresponding message, which will then be published to any potential subscribers.

To prevent orphaned words (i.e, words that have failed to close correctly), any word that is open and that has not had a symbol appended for longer than a defined expiry duration is automatically closed and discarded.

4 Experimental Setup

This section describes the experiments undertaken in order to validate the proposed audio-based communication system. These experiments involve the use of humanoid Nao robots and a purpose built vehicle capable of transporting a Nao robot.

The following subsections describe the process for selecting the length of emitted DTMF audio tones, as well as the tests conducted in order to validate the system.

4.1 DTMF Tone Length

According to [International Telecommunication Union, 1990], the minimum listening time for detecting a single DTMF tone is 40ms. However for the purpose of transmitting tones over an extended distance, the minimum listening time for a single tone was determined empirically to be 80ms.

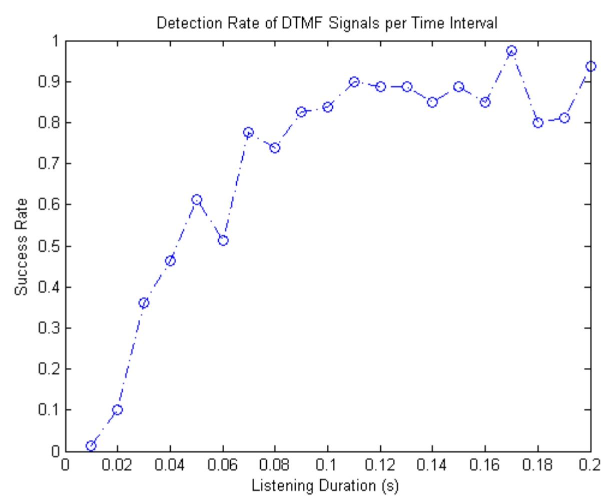

Figure 4: The tone detection rate when the robots were placed at a distance of 1 metre apart.

In order to increase the transmission rate of the system, while ensuring that the reliability of the transmission, a range of transmission times were tested. This range consisted of times between 10ms and 200ms with increments of 10ms between each time. For each time increment, the speaker would transmit each of the 16 available tones 5 times, totalling 80 tones per time increment. Each tone is followed by a 100ms delay in order to prevent tones the receiver inadvertently mixing tones.

The results at each time increment were recorded, and the success rate per time increment calculated. Based on these results, our findings have shown that for a tone detection time of 80ms, the highest reliability was achieved with a tone length of 170ms, with a detection rate of 97.5% (see Figure 4). Using this figure, with the added delay of 100ms, the total transmission time for each tone is 270ms, giving a transmission rate of 3.7 tones per second.

4.2 Tests

Word Detection Robustness

This task is designed to evaluate the reliability of word transmission within an office environment. This task involve transmitting a series of randomly selected words of various symbol lengths between two Nao robots, placed first 1 metre and then 5 metres apart.

Both robots are provided with a dictionary consisting of every possible word between 1 and 3 tones in length for a total of 2562 words. This total excludes invalid words containing two instances of the same symbol in sequence.

Initially placed at a distance of 1 metre apart, the speaking robot is required to randomly select and transmit 50 words of 1 tone in length, 50 words of 2 tones in length, and 50 words of 3 tones in length.

At the end of each transmission, the speaker and listener will record the results of the exchange. These logs include the name of the transmitted message by the speaker, whether an acknowledgement was received, and the name of the message derived by the listener from decoding the transmitted word.

To determine the effect of distance on the reliability of transmission, the robots are then moved to a distance of 5 metres apart, and the experiment repeated.

Follow the Leader Task

This task is intended to demonstrate the ability of the proposed audio-based communication system to facilitate the initiation and termination of a simple follow task between two Nao robots (see Figure 7).

Operating with an office environment, both robots are provided with a dictionary of words containing the words needed to complete this task. This dictionary includes words that have been mapped to messages such as "follow me", "I'm ready" and "we're here".

Initially, both robots are placed several metres apart, facing in opposing directions. Each robot begins the task in an idle state in a crouched pose. One robot will then be instructed (through a touch to the head) to begin the task, acting in the role of leader. The leader will then navigate to, and through the use of the appropriate word, activate the follower role in the second robot. This activation will be trigged by the use of the word representing "Follow Me". The follower will then move into a standing position, and turn in place until it has located the leader.

Once the leader has been located, the follower will then transmit the word representing "I'm ready". The leader will then turn in place, and proceed to walk in the direction of its starting position. The follower is at this stage required to keep pace with the leader at all times until instructed otherwise.

Once the leader has moved a sufficient distance, it will transmit the word representing "we're here" to the follower, terminating the task. Both robots will then return to a crouched pose.

Turn Taking Task - Driving a Vehicle

This task is intended to demonstrate the ability of the proposed audio-based communication system to facilitate a turn-taking task between a humanoid Nao robot and the simple robotic vehicle. This task will involve the Nao robot mounting, dismounting and driving the robotic vehicle (see Figure 8).

Operating within an office environment, both robots are provided with a dictionary containing the words needed to complete the task. This dictionary includes words that have been mapped to messages such as "Dock with me", "I'm ready", "Drive forward" and "Stop".

Initially, the Nao is placed at a distance of 30 centimetres in front of the vehicle with its back to the vehicle. Both the Nao and the vehicle begin the task in an idle state, with the Nao in a crouching pose. On commencement of the task, the Nao moves into a standing position, and indicates to the vehicle that the vehicle move into position against Nao.

Once the vehicle has moved into position, it then communicates to the Nao that it is ready for the Nao to mount. The Nao will then proceed to execute the mounting manoeuvre, at the end of which it will transmit to the vehicle that it is ready to begin driving. The Nao will then instruct the vehicle to move forward a short distance, before instructing it to stop.

Upon stopping, the Nao will dismount from the vehicle, coming to a standing pose without the mounting bracket of the vehicle. The Nao will then inform the vehicle that the task has concluded, causing the vehicle to reverse a safe distance away from the Nao. The vehicle will then instruct the Nao that it is now free to move independently. Both robots then return to an idle state, with the Nao moving into a crouching pose.

5 Experimental Results

Word Detection Robustness

When placed at a distance of 1 metre apart, all single-tone and two-tone long words were successfully detected by the receiver, while 1 three-tone word was missed (see Figure 5). Of the 50 single-tone words detected, 2 words were incorrectly decoded, meaning that the listener detected a different word from what was transmitted. There were no incorrectly decoded words for the two-tone and three-tone words. This provides an overall word detection rate at a distance of 1 metre of 97%.

When the robots were moved to a distance of 5 metres apart, the detection rate was remained unchanged for each of the three groups (see Figure 6). However, the rate of incorrectly decoded words increased slightly. For the single tone-words, 1 incorrect classification was made, for the two-tone words, 5 incorrect classifications were made, and for the three tone words, 1 incorrect classification was made. This provides an overall word detection rate at a distance of 1 metre of 94.6%.

Based on these findings, it can be seen that distance between the speaker and listener is an important factor in the ability of the system to correctly decode transmitted words. However, this could be overcome having the listening robot request confirmation if a decoded word is unexpected or is not relevant to a given situation.

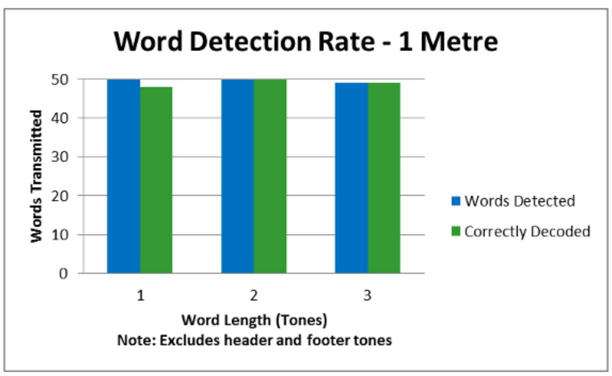

Figure 5: The word detection rate when the robots were placed at a distance of 1 metre apart.

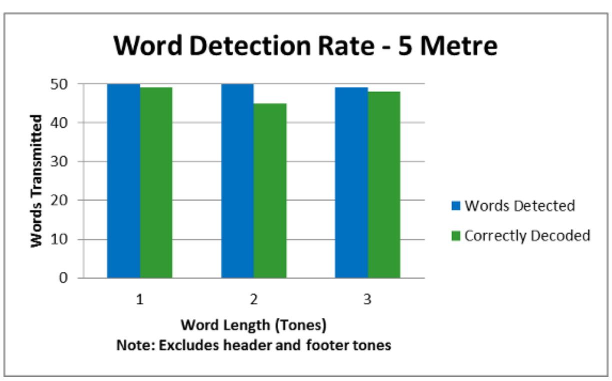

Figure 6: The word detection rate when the robots were placed at a distance of 5 metres apart.

Follow the Leader Task

The experiment was conducted a total of 5 times, in each case under the same conditions. On all 5 occasions, the leader and follower were successfully able to complete the follow task using only audio communication.

Figure 7 shows the progression of the follow task. In Figure 7(a) the lead Nao transmits a request for the follower to begin following, with the follower providing acknowledgement of the request. The result of this exchange can be seen in Figure 7(b), in which the follower is moving in to a standing pose so that it may follow. Figure 7(c) shows the follower following closely behind the leader. Finally, Figure 7(d) shows both robots in an idle state after the leader has indicated the completion of the task.

Turn Taking Task - Driving a Vehicle

The experiment was completed a total of 5 times. On each occasion, the Nao and vehicle were able to successfully complete the turn taking procedure necessary to permit the robot to dock, drive and undock the vehicle.

Figure 8 demonstrates the progression of the mounting procedure. The Nao, seen standing in front of the vehicle in Figure 8(a), audibly transmits its desire to dock with the vehicle. Figure 8(b) shows the vehicle after having moved into position against the Nao. The vehicle then transmits to the Nao that it is ready for the Nao to dock. In Figure 8(c) The Nao begins the mounting procedure. Finally, Figure 8(d) shows the Nao successfully docked

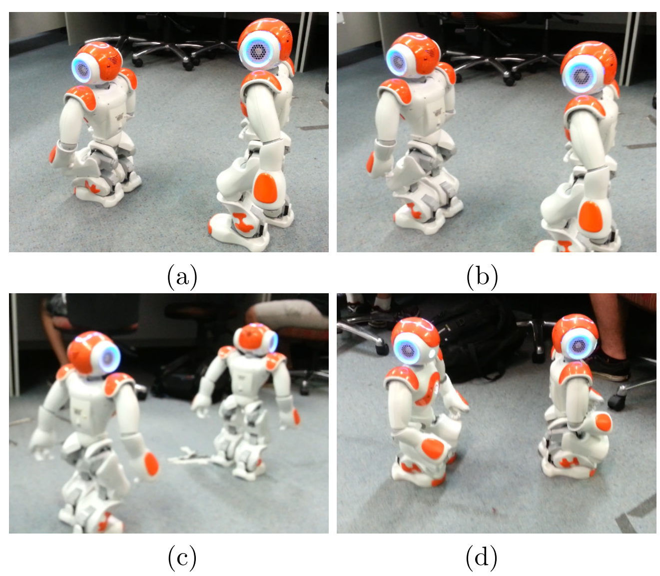

Figure 7: A simple follow task involving two Nao robots communicating through via audio signalling.

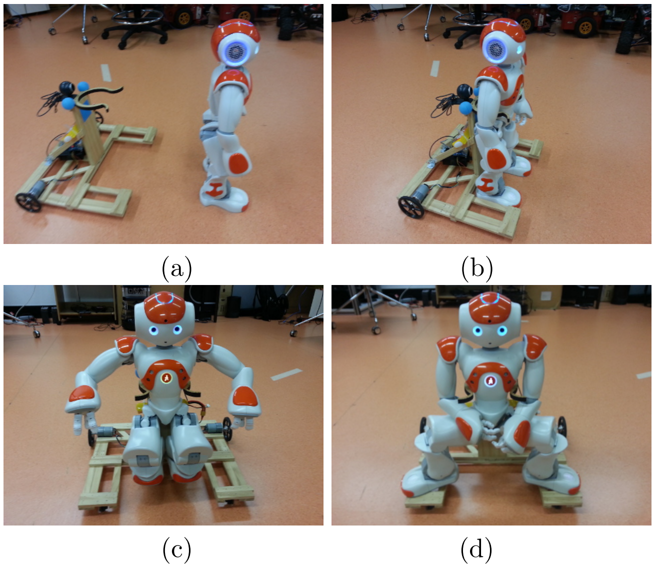

Figure 8: A Nao robot and robotic vehicle performing a turn taking task.

and ready to drive. The Nao then finishes the exchange by transmitting a message to the vehicle that it is now ready to drive.

6 Discussion

The audio communication system presented in this paper is capable of facilitating continued communication between robots that are not capable of accessing a standard network. However, at present, the system does not provide any capabilities analogues to a network link layer. As such, transmissions do not currently contain information that can be used to identify the source or destination of a transmitted word. Therefore, words transmitted by a robot are effectively broadcast to every robot in listening range. This can lead to a response being inadvertently activated in robots that are not the intended recipient of the transmission.

While this situation can potentially be addressed in multi-robot systems containing larger numbers of robots by designating a speaker to handle a defective robot, the communication system would still be incapable of handling a complete network outage. Therefore, in order to allow robots to ascertain the source and destination of transmitted messages, we are currently investigating reserving additional symbols as potential source-destination flags.

Finally, as this system is designed to act as a backup communication channel, robots equipped with the system must have a policy in place for switching to audio communication in the event of a network device failure. Additionally, each robot in the network must be capable of detecting when a robot disappears from the network, either by occasionally polling each robot on the network, or by continually listening for potential audio transmissions.

7 Conclusion

The system presented in this paper demonstrates a novel backup communication system to permit low-bandwidth communication between pairs of robots through the use of audio signalling. By providing access to low-bandwidth communication, our aim has been to demonstrate that robots experiencing a network failure may still participate in a multi-robot system.

Through combining DTMF tones into unique sequences, or words, we were successfully able to demonstrate pairs of robots negotiating a set of tasks. These tasks involved both a simple follow task, as well as a more complex turn taking task involving docking with and driving a robotic vehicle.

Further, our tests have shown that at a transmission rate of 3.7 symbols per second, the system is capable of achieving a word detection reliability rate of 97% at a distance of 1 metre, and 94.6% at a distance of 5 metres.

References

[Basu and Redi, 2004] P. Basu and J. Redi. Movement control algorithms for realization of fault-tolerant ad hoc robot networks. *Network, IEEE*, 18(4):36–44, July 2004.

[Conesa-Munoz *et al.*, 2012] J. Conesa-Munoz, A Ribeiro, D. Andujar, C. Fernandez-Quintanilla, and J. Dorado. Multi-path planning based on a nsga-ii for a fleet of robots to work on agricultural tasks. In *Evolutionary Computation (CEC), 2012 IEEE Congress on*, pages 1–8, June 2012.

[Couceiro *et al.*, 2011] M.S. Couceiro, R.P. Rocha, and N.M.F. Ferreira. Ensuring ad hoc connectivity in distributed search with robotic darwinian particle swarms. In *Safety, Security, and Rescue Robotics (SSRR), 2011 IEEE International Symposium on*, pages 284–289, Nov 2011.

[Couceiro *et al.*, 2013] Micael S. Couceiro, David Portugal, and Rui P. Rocha. A collective robotic architecture in search and rescue scenarios. In *Proceedings of the 28th Annual ACM Symposium on Applied Computing*, SAC '13, pages 64–69, New York, NY, USA, 2013. ACM.

[Eoh *et al.*, 2013] Gyuho Eoh, Jeong S. Choi, and Beom H. Lee. Faulty robot rescue by multi-robot cooperation. *Robotica*, 31:1239–1249, 12 2013.

[Fang and Parker, 2005] T. Fang and L.E. Parker. Distributed multi-robot coalitions through AsyMTRe-D.

In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2606–2613, n.p, 2005.

[Goertzel, 1958] Gerald Goertzel. An algorithm for the evaluation of finite trigonometric series. *The American Mathematical Monthly*, 65(1):pp. 34–35, 1958.

[International Telecommunication Union, 1990] International Telecommunication Union. Q.24: Multifrequency push-button signal reception. https://www.itu.int/rec/T-REC-Q.24-198811-I/en, May 1990.

[Karimian *et al.*, 2006] P. Karimian, R. Vaughan, and S. Brown. Sounds good: Simulation and evaluation of audio communication for multi-robot exploration. In *Intelligent Robots and Systems, 2006 IEEE/RSJ International Conference on*, pages 2711–2716, Oct 2006.

[Parker and Fang, 2006] L.E. Parker and T. Fang. Building multirobot coalitions through automated task solution synthesis. *Proceedings of the IEEE*, 94(7):1289–1305, 2006.

[Parker, 1998] L.E. Parker. ALLIANCE: an architecture for fault tolerant multirobot cooperation. *IEEE Transactions on Robotics and Automation*, 14(2):220–240, 1998.

[Rohrmüller *et al.*, 2012] F. Rohrmüller, D. Wollherr, and M. Buss. MuRoCo: A framework for capability- and situation-aware coalition formation in cooperative multi-robot systems. *Journal of Intelligent and Robotic Systems*, 67(3-4):339–370, 2012.

[Schulz *et al.*, 2011] R. Schulz, G. Wyeth, and J. Wiles. Lingodroids: socially grounding place names in privately grounded cognitive maps. *Adaptive Behavior*, 19(6):409–424, 2011.

[Schulz, 2009] Ruth Jennifer Schulz. *Spatial Language for Mobile Robots: The Formation and Generative Grounding of Toponyms*. University of Queensland, 2009.

[Su *et al.*, 2013] Kuo-Lan Su, Yung-Chin Lin, Bo-Yi Li, and Jr-Hung Guo. Speech-based formation control of the multi-robot system. *Artificial Life and Robotics*, 17(3-4):367–372, 2013.

[Vig and Adams, 2006] L. Vig and J.A. Adams. Multi-robot coalition formation. *IEEE Transactions on Robotics*, 22(4):637–649, 2006.