BNAM 2021
Baltic-Nordic Acoustics Meeting
3rd-5th of May, Oslo, Norway

Distributed acoustic acquisition with low-cost embedded systems

Andres Mondragon Clavijo[1] and Guillaume Dutilleux[2]
Acoustics Research Center, Department of Electronic Systems, NTNU, NO-7034, Norway.

Under some circumstances having an acoustic measurement system handling both reproduction and acquisition becomes very challenging or technically impossible (e.g. long-range outdoor sound propagation experiments or when measuring transmission loss between two rooms). A solution to this problem is distributed acquisition in which the measurement system is not composed of a single measuring device but of separate nodes that communicate wirelessly, and which are assigned specific tasks. This paper describes a prototype of such a measurement system for which the main design constraints were cost and portability. The use of the prototype was limited to the measurement of impulse responses with the swept-sine technique since it allows high signal-to-noise ratio while not requiring a tight synchronization between the clocks of the nodes in the system. The quality of the measurements obtained was checked by a reference to a commercial measurement system in a real-world situation. The obtained results suggest that at a lower cost and reduced size, the prototype offers audio quality comparable to that of commercial systems while adding the flexibility associated with distributed acquisition. Nonetheless, node synchronization proved to be a limiting factor in the usability of the system as it was not possible to achieve accurate timing of the beginning of the acquisition below five milliseconds.

1 Introduction

Low-cost acoustic measurement systems are based on the assumption that for most cases only one device is necessary for the acquisition and reproduction of test signals. The issue with this assumption, is that it limits the scope of acoustic measurements to scenarios in which it is easy to have a wired connection between the electro-acoustical transducers and the sampling device.

Consider for instance outdoor sound propagation experiments in which the distance between source and receiver is long enough to make signal integrity over the cables an issue. Or consider scenarios such as the in-situ measurement of the acoustical parameters of a room partition where having a wired connection between the transducers and the sampling device defeats the purpose of the measurement. Consequently, to design a low-cost acoustic measurement system (AMS) that relies on some form of distributed acquisition, is not only an interesting engineering problem, but also an opportunity to analyse previously unstudied measurement scenarios.

The present paper describes the design and implementation of a prototype measurement system with distributed acquisition capabilities. The prototype consists of two wirelessly synchronized nodes with Raspberry Pi computers as processing units. This paper builds upon the work done with the Acoustics Research Center at NTNU for a specialization project and master's thesis titled "Distributed acoustic acquisition with low-cost embedded systems" [1, 2].

The remainder of this paper is organized as follows. In section 2, some interesting measurement scenarios are described to provide some context on the necessity for distributed acoustic acquisition. In section 3, the proposed distributed measurement system is described from an engineering point of view to explain the employed design methodology. Section

[1] amond47@gmail.com, Now working with Texas Instruments as an embedded software engineer.
[2] guillaume.dutilleux@ntnu.no

4 provides an evaluation of the audio quality and acoustic capabilities of the implemented prototype by comparing with a low-cost commercial AMS. An analysis of the obtained results is presented in section 5 to conclude whether the proof of concept was achieved, and to determine the current shortcomings and possible improvements that can be introduced into a new prototype of the measurement system.

2 The need for low-cost distributed acoustic acquisition

As previously mentioned, there are many measurement scenarios in which having distributed acquisition capabilities would either simplify the execution of the measurement, or it would improve the quality of the results. An example of this is the execution of outdoor sound propagation experiments such as the measurement of the acoustic ground impedance in long range. In this case, the long distance between source and receiver becomes a significant challenge, which so far has been handled by manually synchronizing the acquisition to impulsive sources such as propane cannons [3].

Another example is the execution of measurements in difficult environments. Consider for example acoustic measurements inside caves (Figure 1) to study sound propagation [4]. Using distributed acquisition significantly reduces the logistical effort and complexity of the measurement.

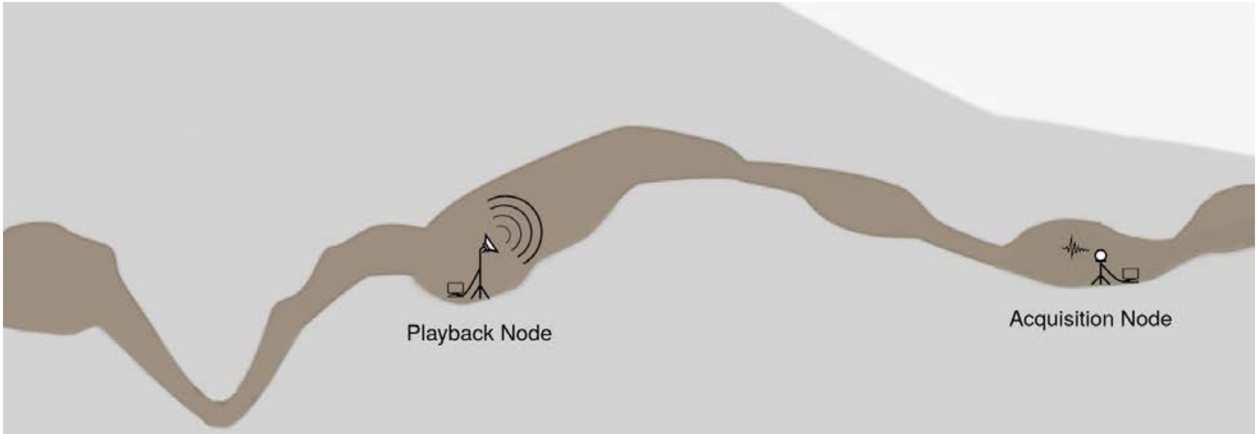

Figure 1: Example scenario of a distributed measurement system (sound propagation experiment inside a cave).

Some manufacturers produce very high-quality measurement systems capable of implementing some form of distributed acquisition. However, a proprietary and closed measurement system might not only be cost-prohibitive for researchers and acousticians, but it might also be inflexible by not allowing them to conceive and develop their own measurements, signal processing and analysis.

These shortcomings of proprietary systems can be solved to some extent by developing a low-cost system that even if it doesn't offer the same level of audio quality as a proprietary system, it does offer flexibility, control, and the capability to conduct distributed acquisition at a low cost.

3 Design and implementation of the Acoustic Measurement System (AMS)

3.1 System's architecture

The design and implementation of the prototype focused on achieving the proof of concept as fast as possible. Therefore, the system mostly relied on commercially available components for hardware, and open source code for software. Selecting a Raspberry Pi single board computer as base for the platform was critical given the vast amount of expansion boards available in the form of HATs (hardware attached on top).

The only component that was specifically designed for the prototype was the analog front-end which was designed to support a 48V phantom power supply for electret condenser microphones and two low-noise microphone pre-amplifiers. Table 1 describes each of the components chosen for each node of the first prototype of the AMS (Figure 2).

Table 1: Software and Hardware components chosen for the prototype AMS.

Component	Design decision	Description
Processing unit	Raspberry Pi	Main processing unit in charge of signal processing and peripheral control.
Audio I/O	HifiBerry DAC+ADC	Includes an audio codec with dedicated 192kHz/24-bit DAC and ADC for a total of two analog inputs and outputs. Characterized by high SNR and low THD and customizable sampling frequency.
Wireless Synchronization	RFM9x LoRa Transceiver	Wireless network interface and transceiver. Relies on LoRa technology to provide wireless long-range communication between nodes.
Signal Conditioning	Analog front-end	Custom made with a 48V phantom power supply and low noise microphone pre-amplifiers.
Signal amplification	HiFiBerry Amp 2	Class-D power amplifier. Replaceable with any other power amplifier depending on requirements for the measurement.
Signal acquisition	Behringer ECM8000	Electret condenser microphone. Replaceable with any other electret condenser microphone that supports phantom power supplies.
Tactile display	Raspberry Pi Display	7" touchscreen used for system control and configuration.
Operating System	Raspbian OS	Operating System in charge of managing hardware and software resources.
Programming Language	Python	In addition to the C language for low level tasks such as audio acquisition. Chosen due to the vast amount of signal processing modules.
User Interface	Tkinter	Graphical User Interface (GUI) to provide the user a way to set measurements and visualize results.
Node synchronization	Timing-sync protocol for sensor Networks (TPSN)	Protocol used for the time synchronization between the nodes of the system.

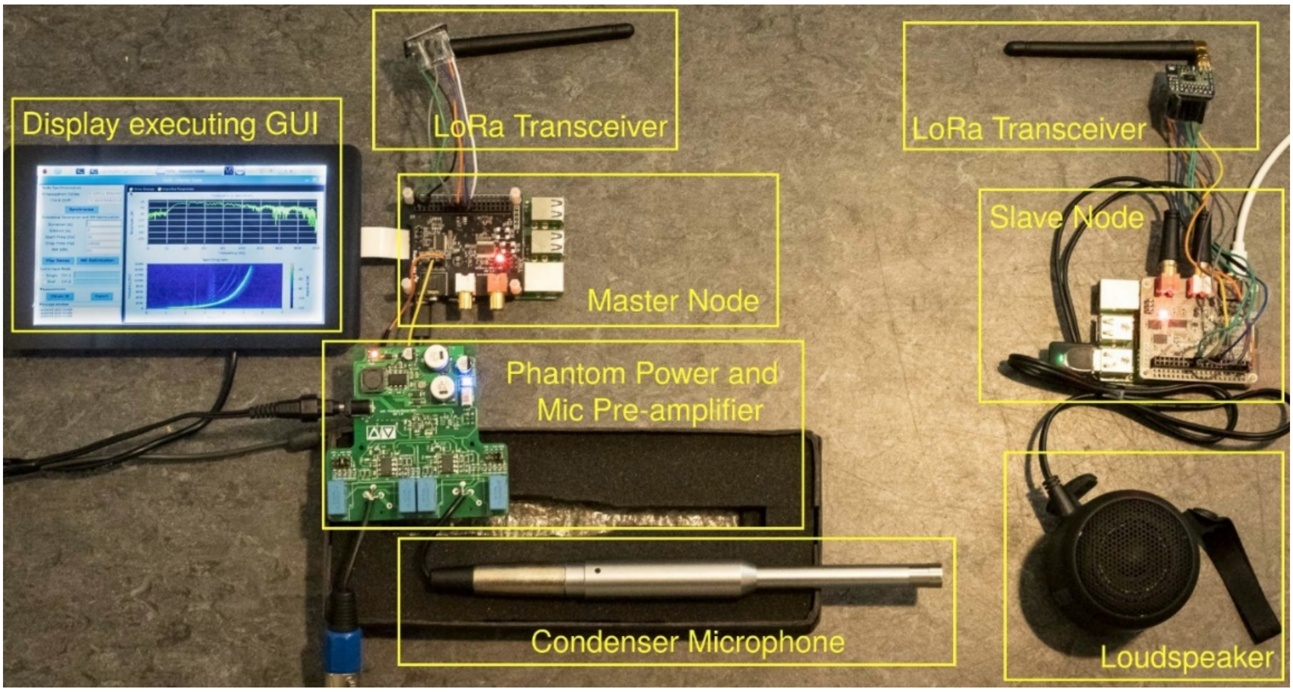

Figure 2: Implemented prototype AMS with two nodes showing all the components (small loudspeaker just for illustrative purposes).

3.2 Impulse response derivation

Early in the design process it was decided that the measurement of impulse responses (IRs) would be one of the functional requirements of the system. This is simply because it is one of the most common tasks in acoustics, and because the measurement of impulse responses and its associated transfer function allows acousticians to analyze a wide variety of acoustical and electro-acoustical systems. Taking this into account, determining which measurement technique should be used to measure impulse responses was a critical part of the system's design. Here it is important not only to consider aspects such as high SNR and noise rejection, but also aspects related to the distributed acquisition capabilities of the system. For instance, pseudo-random techniques such as the one based on the use of maximum length sequences (MLS) is not compatible with distributed acquisition because it requires a very tight synchronization between the sampler used to reproduce the signal and the one used to record the system's response [5, 6].

Of the different techniques considered, the swept-sine method proved to be the most suitable for this particular application [7]. This is mainly because of its high SNR, impulsive noise immunity, non-susceptibility to harmonic distortion products, and most important due to the fact that there is no need to maintain a tight synchronization between sampling clock of the signal generator and the one of the digitizing unit employed for capturing the system's response [7, 8, 9].

3.3 Node synchronization and prototype implementation

In this particular context, node synchronization can be understood in two ways. A strict definition implies that the distributed system must be capable of synchronizing every playback sample in one node to an acquisition sample in the other node, something which is extremely challenging. A less strict definition of node synchronization considers only two requirements. First, there shouldn't be a clock mismatch between the nodes, and second, the time difference between the start of playback and the start of acquisition should be small enough not to have an effect on the measurement.

Node synchronization in the prototype relied on a simplified implementation of a synchronization protocol known as the Timing-sync Protocol for Sensor Networks (TPSN) [10]. The critical aspect of the synchronization is the pairwise synchronization between the nodes, which uses a sender-receiver synchronization approach. By sending synchronization pulses with specific timestamps, it is possible to calculate the clock drift and the propagation delay. These values can be then used to synchronize the nodes as often as needed.

4 Evaluation of the implemented prototype

After an initial assessment of the prototype's overall operation, and a verification of the full audio chain, a series of detailed tests took place to evaluate audio quality, distributed acquisition, and impulse response derivation. To provide a reference point for the analysis, these tests were also executed (when applicable) on a low-cost measurement system composed of a laptop computer, an Edirol UA-25 consumer-grade sound card, and the EASERA acoustics software. For a more detailed account of the measurements and obtained results refer to [1, 2].

4.1 Distributed operation and impulse response derivation

Distributed acquisition was evaluated in two distinct manners. The first one consisted in sequentially executing a series of synchronization attempts to determine synchronization accuracy and precision. This was done to offer some insights on the type of acoustic measurements that could be handled by the prototype. The second one consisted in performing a controlled closed-loop distributed measurement to obtain and analyse the system's own impulse response. The measurement involves connecting the output of the slave node to the input of the master node and performing a distributed acquisition (distributed in the sense that wireless synchronization is used to synchronize when playback starts in one node, and acquisition starts in the other node). This type of measurement allows to characterize the system itself by providing information about the system's frequency response, harmonic distortion presence, and clock mismatch between the nodes.

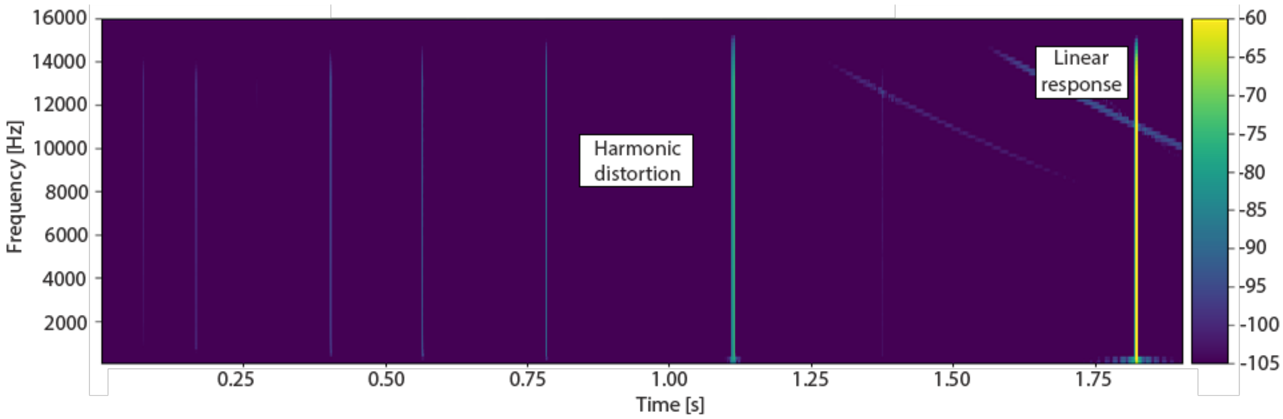

Figure 3: Obtained impulse response for a closed-loop measurement (sine sweep limited to the 31 Hz – 16 kHz frequency range).

A good measurement system can be described as one that is invisible to the measurement, that is, that the results of the measurement are not influenced by the systems' own response. The frequency response of the system should be as flat as possible, and the non-linear behaviour of the system shouldn't affect the measurement results. These two properties can be evaluated by visualizing the system's response in the time and in the frequency domains.

As seen in Figure 3, harmonic distortion associated with the non-linear behaviour of the system is kept under control by using the swept-sine technique. In fact, this is one of the major advantages of the technique. Since the system's linear response is separated from the non-linear response in the time domain, it is easy to discard the non-linearities of the system.

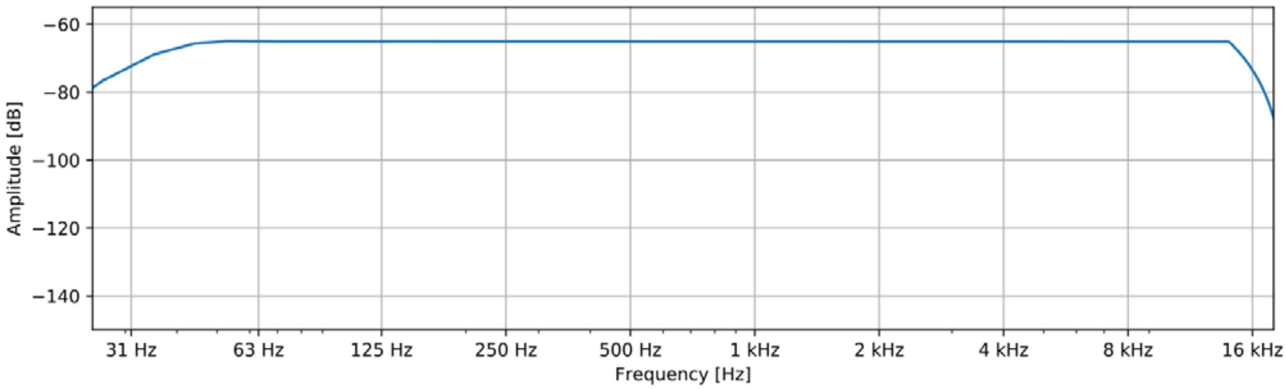

Figure 4: Obtained linear frequency response for a closed-loop measurement (sine sweep limited to the 31 Hz – 16 kHz frequency range).

On the other hand, a flat frequency response indicates that all the components in the audio chain are transparent to the measurement, indicating that the audio codec won't affect the quality of the measurement (Figure 4).

One of the main issues associated with distributed acquisition is the possibility of a clock mismatch between the nodes. This occurs because each node has a dedicated clock to provide accurate timing for the sampling process and depending on the chosen measurement technique, it can significantly affect impulse response derivation.

The swept-sine technique is particularly robust against a clock mismatch, which is the main reason for choosing it for a distributed acquisition system. Nevertheless, if there is a clock mismatch between the nodes, the resulting impulse response will display what is known as skewness (a deformation of the impulse response) that can be visualized with a spectrogram.

Even though it wasn't possible to observe a clock mismatch between the nodes when conducting the closed-loop measurement (even with very long sine sweeps), it was possible to conduct a measurement in which the sampling rate of one of the nodes was slightly changed to produce a clock mismatch (Figure 5). This was possible due to the system's ability to control the sampling rate of the nodes, that is, to precisely configure sampling rates different than the standard ones such as 44.1kHz, 48 kHz, 192kHz, etc.

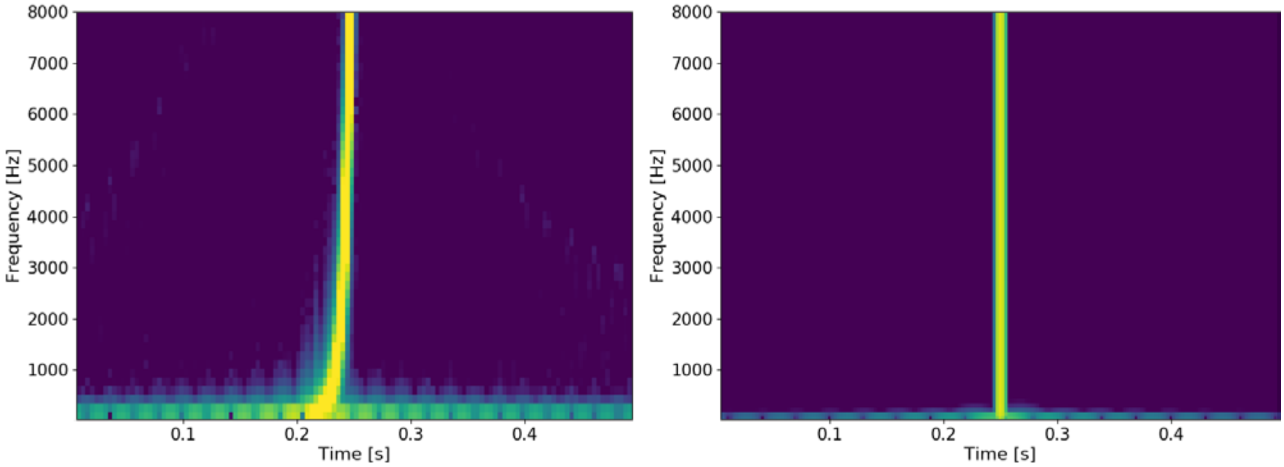

Figure 5: Simulated skewed IR with a clock mismatch of 0.1% (left) and real IR obtained with a closed-loop measurement.

Finally, the precision of the node synchronization was established by repeatedly executing the pairwise synchronization (according to the TPSN) of the nodes and determining the variation between the obtained values for propagation delay. This measurement allowed to calculate a standard deviation of 7.14 milliseconds for the propagation delay in a set of twenty sequential synchronization attempts. A value much greater than expected and that can limit the type of measurements that can be executed with the system. Nevertheless, given the use of the swept-sine technique, the effect of this lack of precision is minor and can be significantly minimized by using longer sine sweeps.

4.2 Test measurement scenario (Reverberation Time)

Taking into consideration that the prototype is meant to be used in real measurement scenarios, a series of reverberation time measurements took place in one of the reverberation chambers located in the acoustics lab at NTNU. The measurement was done such that time-invariance due to environmental factors (e.g. temperature) would be reduced, and proper energy distribution would be ensured by a set of loudspeakers located at specific points in the room.

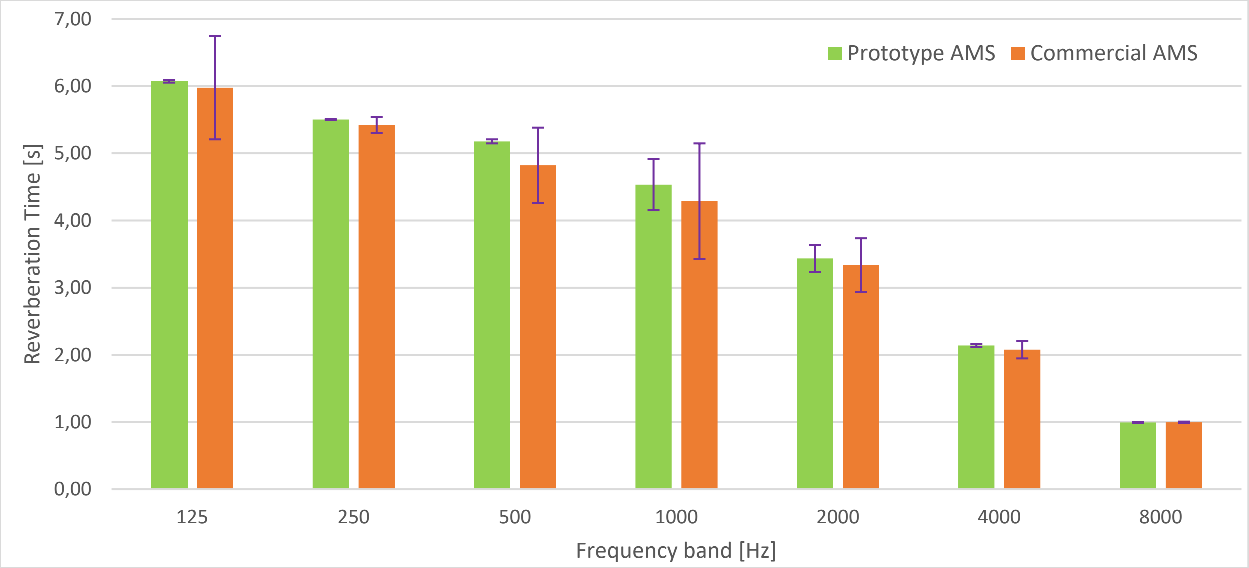

Figure 6: Calculation of reverberation time with the EDT descriptor in EASERA. The error bars correspond to +/- 1σ.

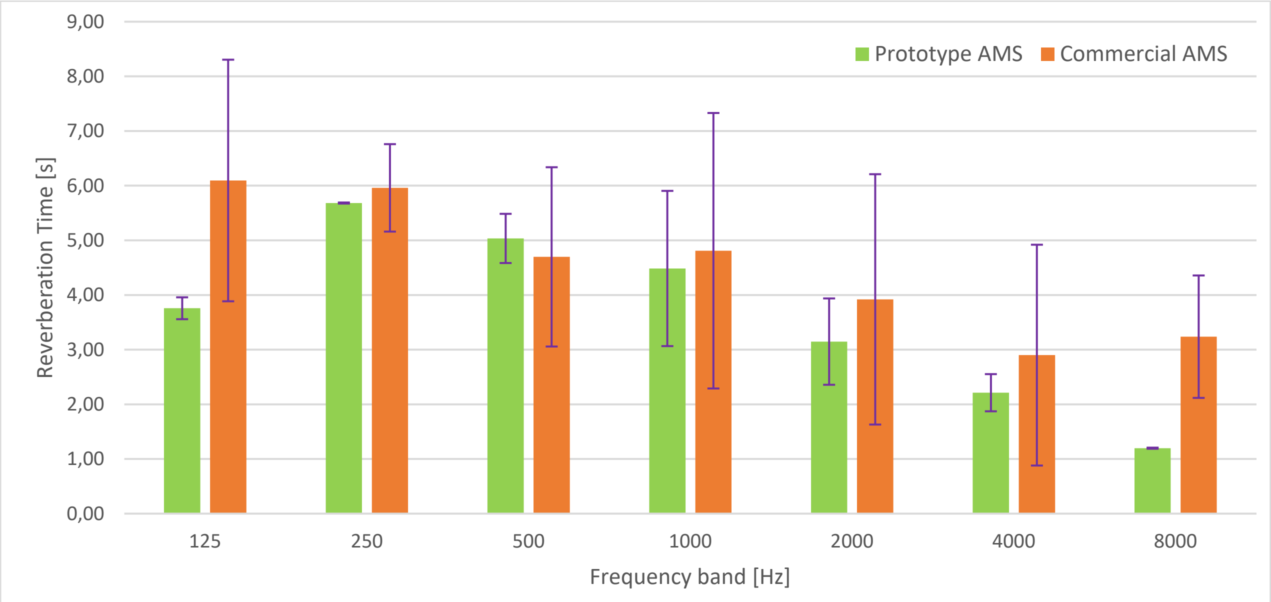

Figure 7: Calculation of reverberation time with the T30 descriptor in EASERA. The error bars correspond to +/- 1σ.

Four IR measurements were consecutively performed to obtain the IRs with both the prototype AMS and the low-cost commercial AMS (based on EASERA). Once the IRs were obtained and saved as audio files with each measurement system, EASERA was used as the analysis tool to derive the reverberation time from the EDT and the T30 acoustic descriptors. With a reverberation time for each measurement, the standard deviation for the set of measurements was calculated and then used to determine how precise each measurement system was (Figure 6 and Figure 7).

5 Analysis of the obtained results

The closed-loop measurement executed on the system suggests that the prototype system can perform acoustic measurements without having a significant influence on the measurement results. The prototype AMS has a flat frequency response throughout most of its frequency range and harmonic distortion artifacts can be completely removed from the measurement by using the swept-sine technique. Nevertheless, Figure 4 suggests that the start frequency and the stop frequency of the sine sweep must be adjusted so that the entirety of the desired frequency range is covered. Even though the sine sweep used for the measurement was supposed to cover the entire range from 31 Hz to 16 kHz, the obtained IR suggests otherwise. This might be a consequence of the sine sweep processing that was employed to minimize problems associated with pre-ringing by using fade-ins and fade-outs.

To determine the presence of a clock mismatch, long sine sweeps, with a duration of sixty seconds were used as test signals. The obtained impulse responses were then analysed with a spectrogram and showed no sign of skewedness, which was validated by artificially creating a clock mismatch between the nodes. And even though no clock mismatch was detected, the system's ability to accurately adjust the sampling frequency of its nodes indicates that the system could adapt to circumstances in which a clock mismatch does occur, perhaps due to environmental conditions such as humidity or temperature differences between the nodes.

The results for node synchronization were disappointing given that the TPSN protocol can theoretically synchronize nodes with an average accuracy of less than 20µs. It was determined that the use of a non-deterministic operating system (i.e. Raspbian) significantly affects the level of synchronization and that better results could be obtained using either bare-metal programming or a real-time operating system in which tasks can be fully executed within a given time deadline.

Nonetheless, the results also suggest that the lack of precise node synchronization does not significantly affect the execution of acoustic measurements. Both the closed-loop measurements performed to evaluate the acoustic capabilities of the system, and the preliminary reverberation time measurements conducted at NTNU suggest that the system is suitable for acoustic measurements.

In fact, results also suggest that the prototype produces more precise measurements than the low-cost commercial AMS used for comparison. The standard deviation was significantly lower in the measurements executed with the prototype (see Figure 6 and Figure 7). This could be due to different reasons, but one hypothesis is that the swept-sine formulation used in the prototype (formulated in [7] and further refined in [9]) is more robust than the one used in EASERA. Another likely hypothesis is that a sub-optimal setting of the sound card used with the commercial system affected the quality of the measurements. More measurements in different scenarios are required to provide more conclusive results.

6 Conclusions and final remarks

We have developed a prototype of a portable 2-channel low-cost distributed acoustic IR measurement system with phantom-powering. The system consists of two nodes. Each node can be used independently in full-duplex mode, or as part of a distributed system with one master input node that wirelessly triggers the playback of the test signal from a slave node. The IR acquisition relies on exponential sine sweeps.

The main objective of the prototype was to evaluate the viability of this system. In this sense, it is possible to say that the proof of concept was achieved, and that future work should focus on improving audio quality and making the system more robust and streamlined. This means, redesigning the system with more appropriate components such as a processing unit with less software overhead. The implementation of the analog front-end although successful requires more work to fully take advantage of the high SNR and low THD of the audio codec.

Furthermore, it is necessary to conduct additional measurements to determine the range of the radio transmitters under different conditions (with and without line of sight). Moreover, the proposed prototype is not limited to only two nodes and can be further extended (in practical terms) to as many nodes as the user might need. This would require some minor modifications to the synchronization between nodes, but there is no technical challenge associated with this.

Finally, even though the measurement results presented in this paper can be qualified only as preliminary, they are good enough to plan more objective comparative measurements with real equipment and in more interesting measurement scenarios.

References

[1] A. M. Clavijo, Distributed acoustic acquisition with low-cost embedded systems, *Specialization Project*, Department of Electronic Systems (Unpublished Work), NTNU, 2019.

[2] A. M. Clavijo, Distributed acoustic acquisition with low-cost embedded systems, *Master's thesis*, European Master in Embedded Computing Systems (EMECS), NTNU, 2020.

[3] J. H. Soh, K. E. Gilbert, W. G. Frazier, C. L. Talmadge, and R. Waxler, A direct method for measuring acoustic ground impedance in long-range propagation experiments, *The Journal of the Acoustical Society of America*, vol. 128, no. 5, pp. EL286–EL293, 2010.

[4] N. Yioutsos, G. Kamaris, K. Kaleris, C. Papadakos, J. Mourjopoulos, Archaeoacoustic Research on Caves dedicated to Pan and the Nymphs in Attica, Greece, *Euronoise 2018 Conference*, Crete, 2018.

[5] S. Müller, and P. Massarani, Transfer-function measurement with sweeps, *J. Audio Eng. Soc*, vol. 49, no. 6, pp. 443–471, 2001.

[6] J. S. Abel, N. J. Bryan, and M. A. Kolar, Impulse response measurements in the presence of clock drift, *Audio Engineering Society Convention 129*, Nov 2010.

[7] A. Farina, Simultaneous measurement of impulse response and distortion with a swept-sine technique, *Audio Engineering Society Convention 108*, Feb 2000.

[8] A. Farina, Advancements in impulse response measurements by sine sweeps, *Audio Engineering Society Convention 122*, May 2007.

[9] A. Novak, Identification of Nonlinear Systems in Acoustics. *PhD thesis*, Universite du Maine, 2009.

[10] S. Ganeriwal, R. Kumar, and M. B. Srivastava, Timing-sync protocol for sensor networks, *SenSys 2003: Proceedings of the 1st international conference on Embedded networked sensor systems*, pp. 138–149, Nov 2003.