

Document downloaded from:

<http://hdl.handle.net/10251/188947>

This paper must be cited as:

Belloch, JA.; Badía, JM.; Larios, DF.; Personal, E.; Ferrer Contreras, M.; Fuster Criado, L.; Lupoiu, M.... (2021). On the performance of a GPU-based SoC in a distributed spatial audio system. *The Journal of Supercomputing* (Online). 77(7):6920-6935.  
<https://doi.org/10.1007/s11227-020-03577-4>



The final publication is available at

<https://doi.org/10.1007/s11227-020-03577-4>

Copyright Springer-Verlag

Additional Information

Noname manuscript No.  
(will be inserted by the editor)

---

## On the Performance of a GPU-based SoC in a distributed Spatial Audio System

Jose A. Belloch · José M. Badía ·  
Diego Larios · Enrique Personal ·  
Miguel Ferrer · Laura Fuster ·  
Mihaita Lupoiu · Alberto Gonzalez ·  
Carlos León · Antonio M. Vidal ·  
Enrique S. Quintana-Ortí

Received: date / Accepted: date

**Abstract** Many current System-on-Chip (SoC) devices are composed of low-power multicore processors combined with a small graphics accelerator (or GPU) offering a trade-off between computational capacity and low-power consumption. In this context, spatial audio methods such as Wave Field Synthesis (WFS) can benefit from a distributed system composed of several SoCs that collaborate to tackle the high computational cost of rendering virtual sound sources. This paper aims at evaluating important aspects dealing with a distributed WFS implementation that runs over a network of Jetson Nano boards composed of embedded GPU-based SoCs: computational performance, energy efficiency, and synchronization issues. Our results show that the maximum efficiency is obtained when the WFS system operates the GPU frequency at 691.2 MHz, achieving 11 sources-per-Watt. Synchronization experiments using the NTP protocol show that the maximum initial delay of 10 milliseconds between nodes does not prevent us from achieving high spatial sound quality.

**Keywords** wave field synthesis, spatial audio, real time, embedded systems, GPU, Jetson Nano, System-on-Chip (SoC)

---

Jose A. Belloch  
Depto. de Tecnología Electrónica, Universidad Carlos III de Madrid, Spain  
E-mail: jbelloc@ing.uc3m.es

José M. Badía  
Depto. de Ingeniería y Ciencia de Computadores, Universitat Jaume I de Castellón, Spain.  
E-mail: badia@uji.es

Diego Larios, Enrique Personal and Carlos León  
Depto. de Tecnología Electrónica, Universidad de Sevilla, Spain  
E-mail: {dlarios,eperonal,cleon}@us.es

Alberto Gonzalez, Miguel Ferrer, Antonio M. Vidal,  
Enrique S. Quintana-Ortí, Laura Fuster and Mihaita Lupoiu  
Universitat Politècnica de València, Spain  
E-mail: {agonzal,mferrer}@dcom.upv.es, avidal@dsic.upv.es, quintana@disca.upv.es,  
lfuster@iteam.upv.es, myhay.360@gmail.com

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65

## 1 Introduction

In the last decade, parallel systems are being employed in all segments of the industry in the form of multicore processors and many-core hardware accelerators. Digital signal processing is one of the fields which has largely benefited from these devices, and more specifically the synthesis of the spatial audio.

One of the spatial audio technologies available today is Wave Field Synthesis (WFS) in which a sound field is synthesized in a wide area by means of loudspeaker arrays. WFS is usually tackled via digital signal processing techniques to reproduce complex auditory scenes consisting of multiple acoustic objects. The WFS concept was introduced at the Delft University of Technology in the 1980s by Berkhout [1].

A large-scale WFS system requires performing costly computational operations in real time since it involves multiples input and output channels. In recent years, several efforts have aimed at implementing efficiently a WFS system [2]. Some of the authors of the present work carried out an implementation that enhanced the audio quality by offloading the additional computational cost to Graphics Processing Units (GPUs) [3]. To this end, in addition to rendering virtual sound sources, the aforementioned implementation took into account the computational cost that appears when a Room Compensation (RC) block is added to the system. Common RC blocks are based on multichannel inverse filter banks and their purpose is to correct the room effects or echoes at selected points within the listening area [4, 5]. The proposed system provided high performance in terms of computational and acoustic aspects [3]. However, this WFS implementation was designed for an specific room that required a centralized system with a dedicated power-hungry GPU. This fact limits the flexibility and scalability of the system, because any change in the number of loudspeakers and/or the hall distribution involves necessarily a reconfiguration of the system.

Nowadays, there exist System-on-Chips (SoCs) that deliver notable computational capacity while partially retaining the appealing low power consumption. This high performance capability is often provided by an embedded graphics accelerator (GPU) that is integrated in the SoC. One example of this type of system is the NVIDIA Jetson Nano board [6] which contains a high-performance 128-core Nvidia Maxwell GPU. One of its main features is its low power consumption combined with several levels of parallelism yielding a very high performance per Watt. Moreover, there exist also the possibility of adjusting its energy consumption by reducing the frequency of the CPU cores or the GPU.

This work introduces a distributed system containing multiple NVIDIA Jetson Nano boards, where each one of them is in charge of managing the output processing signal of a batch of loudspeakers. As the main result from this work, we are build an environment-adaptable and quickly-reconfigurable dynamic WFS system. From the computational point of view, the total cost is shared among the embedded GPUs of the Jetson boards. In this work we analyze the performance of the system in terms of audio processing (virtual

1 sound sources and sizes of RC filters), use of different GPU and CPU operat-  
2 ing frequencies, and energy consumption. The aim of this work is to find an  
3 efficient implementation that exhibits a proper trade-off between performance  
4 and energy consumption taking into account real-time constraints.  
5

6 From the acoustic point of view, each Jetson Nano acts as a node inside  
7 a network of acoustic nodes. All nodes aim to generate collaboratively sound  
8 fields with natural temporal and spatial properties within a volume or area  
9 bounded by secondary sources (arrays of loudspeakers). Therefore, in a dis-  
10 tributed system, the secondary sources are distributed among the different  
11 acoustic nodes. As a consequence, besides considering the computational and  
12 energy aspects, a good synchronization between the nodes is essential in order  
13 to reproduce in unison through the loudspeakers of the system and thus, to  
14 offer a correct spatial audio sensation. Perceptual tests show that the NTP  
15 synchronization protocol provides a fair sensation. In order to support this  
16 outcome, we carried out a quantitative experiment using multi-channel audio  
17 cards that were devoted to acquire the audio outputs of all acoustic nodes  
18 (Jetson Nano audio outputs). The proposed experiment allowed us to measure  
19 real output delays among nodes by means of cross-correlations.  
20

21 This paper is structured as follows. Section II offers a brief overview of  
22 distributed algorithms in the field of audio as well as some background in  
23 WFS theory together with a RC filters block. Section III analyzes the node  
24 synchronization in the distributed WFS system. In Section IV, we explore  
25 the performance of the distributed system in terms of maximum number of  
26 sound sources that can be rendered in real time; provide a detailed analysis  
27 of the power dissipation; and analyze the energy efficiency of different hard-  
28 ware configurations. Finally, Section V closes the paper with a few concluding  
29 remarks.  
30

## 31 **2 Audio distributed systems**

34 Distributed systems provide a cheap, flexible and efficient solution for environ-  
35 mental and habitat monitoring, as well as for monitoring and maintenance of  
36 industrial equipment [7]. These systems often involve acoustic applications [8].  
37 Among them, we highlight the use of distributed systems for building huge 3D  
38 spatial audio systems, where the acoustic node is focused on the design of the  
39 output signals. In addition to producing spatial audio effect, the output signals  
40 control the sound field by reducing possible acoustic artifacts that are intro-  
41 duced by external agents. To this end, our designed audio system implements  
42 the WFS technique together with a RC filter block.  
43  
44

### 45 **2.1 Wave Field Synthesis**

46 WFS is a sound rendering method based on the Huygens' principle [9]. The  
47 propagation of a wave front can be described by adding the contribution of a  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65

number of secondary-point sources (arrays of loudspeakers) distributed along the wave front, as shown in Fig. 1(a). Following a model-based rendering with point sources and plane waves [10], the field rendered by a sound source  $m$  at a point  $R$ , within the area surrounded by  $N$  loudspeakers, can be expressed as

$$P(\mathbf{x}_R, \omega) = \sum_{n=0}^{N-1} Q_n(\mathbf{x}_m, \omega) \frac{e^{-j\omega r_{nR}/c}}{r_{nR}}, \quad (1)$$

where  $c$  is the speed of the sound,  $\mathbf{x}_m$  is the position of the virtual sound source  $m$ ,  $\mathbf{x}_R$  is the position of the point  $R$ , and  $r_{nR} = |\mathbf{x}_n - \mathbf{x}_R|$  is the distance between the  $n$ -th loudspeaker and the point  $R$ .

The signal to be reproduced by the  $n$ -th loudspeaker is represented by  $Q_n(\mathbf{x}_m, \omega)$ , which is given by


$$Q_n(\mathbf{x}_m, \omega) = S(\omega) \sqrt{\frac{j\omega}{2\pi c}} K \frac{1}{\sqrt{r_{mn}}} \cos(\theta_{mn}) e^{-j\omega r_{mn}/c}, \quad (2)$$

where  $K$  is a geometry-dependent constant,  $r_{mn} = |\mathbf{x}_m - \mathbf{x}_n|$  and  $\mathbf{x}_n$  is the position of the loudspeaker  $n$ . Fig. 1(b) shows the geometry of the system, where  $\theta_{mn}$  denotes the angle between the line that connects  $\mathbf{x}_m$  and  $\mathbf{x}_n$  and the normal vector  $\mathbf{n}$  of the loudspeaker  $n$ . The guitar represents the sound source  $m$ .

The computed signal (2) consists of several elements that aim to add spatial audio properties to the original source signal whose frequency-domain characteristic is represented by the term  $S(\omega)$ . More details on this technique can be found in [10].

## 2.2 Room Compensation Filters

The synthesized sound field can be altered with new echoes introduced by the listening room, reducing the spatial effect. In [5, 11], the authors designed and validated a multichannel inverse filter bank that corrects these room effects at



**Fig. 1** (a) Theoretical basis for WFS that are based on Huygens' principle. (b) Geometry of a WFS system with the sound source  $m$  (guitar),  $N$  loudspeakers, and the distances among the sound source, the loudspeakers, and the listener ( $R$ ).

1 selected points within the listening area. However, in a WFS system composed  
2 of  $N$  loudspeakers, this implies inserting  $N^2$  FIR filters to the system, con-  
3 siderably increasing its computational demand. It is important to highlight  
4 that the length of the filters can be customized before the execution of the  
5 WFS distributed system so that we can select higher lengths, which involves  
6 a more accurate synthesized signal. In this sense, large lengths avoid that the  
7 generated spatial sound is affected by external aspects such as room effects or  
8 loudspeakers setup, and provide high spatial sound quality [5]. Unfortunately,  
9 this also constrains the number of sound sources that can be rendered in real  
10 time.  
11

### 12 3 Distributed system set-up and synchronization issues

13  
14 In order to tackle synchronization issues, we set a small distributed WFS  
15 system as a testbench (see Fig. 2(a)) composed of the following components:

16 **WFS Master:** Its role is to send the audio samples to all WFS nodes.

17 **Wireless Access Point:** It is used for connecting the different network ele-  
18 ments.


19 **Reference Time Server:** It broadcasts the reference clock to all WFS nodes.

20 **Audio nodes:** They are devoted to synthesize and reproduce the audio through-  
21 out their corresponding loudspeakers.  
22

23 Figure 2(b) shows a detail of a WFS node based on a NVIDIA<sup>®</sup> Jetson  
24 Nano<sup>™</sup> device with an external Wi-Fi and audio USB adapters. Additionally,  
25 Figure 2(c) shows the Reference Time Server, implemented with Raspberry  
26 Pi 3B+. We use an independent device to guarantee that all nodes are in the  
27 same tier level and are affected by the same delays. By choosing a Raspberry  
28 Pi we also reduce the economic cost of the system. Finally, Figure 2(d) shows  
29 three WFS nodes.  
30

31 A Client-Server application based on stream sockets was implemented to  
32 coordinate the different elements of the system. The server allows a dynamic  
33 interaction between the clients, indicating them when to start the reproduction  
34 of one or more sound sources on a given position; when to change the position of  
35 a source; or when to stop its reproduction. All clients read from a configuration  
36 file the number and location of the audio files; the number of loudspeakers and  
37 sound sources; as well as their positions in the system. Each node automatically  
38 computes its own outputs based on equation (2).  
39

40 It is important to note that traditional centralized systems have a com-  
41 mon clock source, so that synchronizing their DACs and audio data signals  
42 is straightforward. Conversely, the distributed strategy proposed in this work  
43 poses the problem that each node has its own clock. Hence, a synchronization  
44 mechanism between the nodes must be enforced. There are different synchron-  
45 ization procedures for networks in the literature, with the most used proto-  
46 col for time synchronization being the Network Time Protocol (NTP) [12].  
47 However, there are more accurate alternatives. The Precision Time Protocol  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65



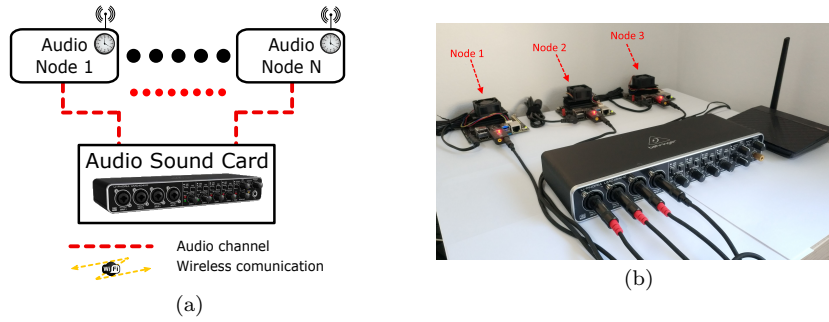
**Fig. 2** Proposed distributed WFS architecture: (a) Interconnection diagram and elements. (b) Acoustic node based on NVIDIA<sup>®</sup> Jetson Nano<sup>™</sup>. (c) NTP server based on Raspberry Pi 3B+. (d) Distributed WFS system prototype; three nodes and six loudspeakers.

(PTP) [13,14] is also a very widespread option; and for applications with very high precision requirements, a synchronization through GPS signal is used [15]. Unfortunately, the last two options require additional hardware support. Also, even though PTP can be emulated via software, its performance gets significantly worse [16]. Due to this, we selected NTP as the synchronization protocol.

We initially carried out perceptual tests when the whole system was settled in order to evaluate the quality of the spatial sound. The people participating in these tests were able to locate easily the origin of the sound source. They highlighted the quality of the sound source movement which was really very smooth. Test results were similar to the tests of the centralized WFS system carried out in [3], so that the audio quality of distributed system resembles that of the centralized system.

### 3.1 Synchronization analysis

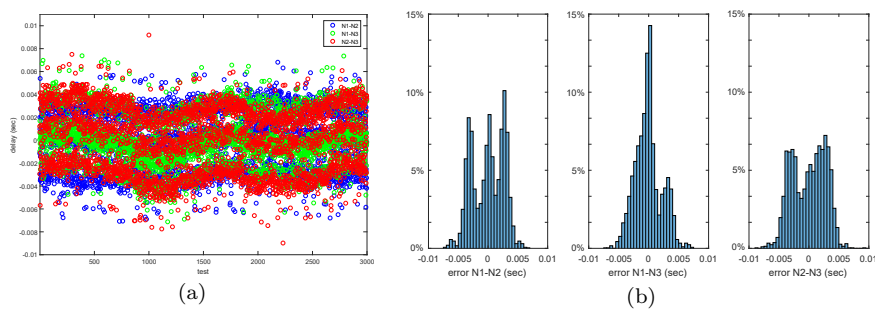
Several tests were carried out to test the usefulness of the proposed system as well as the effects of the synchronization. In these tests, the final elements of the system (loudspeakers) were replaced by a multi-channel audio card, which allows us to acquire the different outputs from the WFS nodes; see Figure 3(a). Specifically, Figure 3(b) shows a small prototype with three nodes and an audio card used for the test.



**Fig. 3** Proposed synchronization test-beds: (a) Audio card connection to the WFS nodes under test. (b) Small prototype set-up with three WFS nodes.

To validate the NTP synchronization method, we configured all the nodes to generate the same audio pattern at their output. Then, we performed a correlation of each captured signal with the pattern signal in order to determine the delay between nodes. Figure 4 shows the delays between the three test-bed nodes for 3,000 test cycles. Each cycle plays a 10 second audio pattern (with a logarithmic sweep track), introducing a random pause of duration between 1 and 5 minutes after it.

As it is shown in Figure 4(a), the delay between nodes in all test was always under 10 milliseconds. Furthermore, by analyzing the dispersion of the results (see Fig. 4(b)), we observe that 97.67%, 98.27% and 97.33% of the



**Fig. 4** Synchronization test results: (a) Recorded delays. (b) Delays distribution.



1 delays between nodes N1-N2, N1-N3 and N2-N3, respectively, were under 5  
2 milliseconds.

3 This quantitative analysis confirms the results obtained with the percep-  
4 tual tests: The NTP protocol provides a precise synchronization method to  
5 implement a distributed WFS system.  
6

## 7 8 9 **4 Computational and energy performance**


10 The NVIDIA Jetson Nano Developer Kit is a board that comprises a SoC  
11 based on a quad-core ARM Cortex-A57 CPU at 1.43 GHz with 4 Gbytes of  
12 LPDDR4 memory, and a 128-core Nvidia Maxwell GPU. One of the main  
13 features of this board is its low-power consumption combined with different  
14 levels of parallelism, which provides a very high performance-per-Watt. A sec-  
15 ond appealing feature of this platform is the possibility of adjusting its energy  
16 consumption by reducing the frequency of the CPU cores or the GPU. The  
17 CPU power management strategy uses dynamic frequency scaling with dy-  
18 namic voltage scaling. The CPU frequency is dynamically adjusted depending  
19 on how busy is the device. However, the GPU and CPU frequencies can be  
20 modified by from inside a user’s code using some features of the Linux kernel.  
21

22 In our experiments, we measured the power dissipation of the Jetson Nano  
23 using the `pmlib` framework [17] to collect the instantaneous power readings  
24 from the internal energy-monitoring sensors [6]. The Jetson Nano board em-  
25 beds real-time current sensors that can be sampled to obtain the power con-  
26 sumption of three separate power domains: main module, GPU and CPU.  
27

### 28 29 30 **4.1 Experimental evaluation**

31 The following experiments evaluate the computational performance, energy  
32 consumption, and energy efficiency of the CUDA-based implementation of  
33 WFS and RC filter block proposed in [3] on one of the Jetson Nano devices  
34 shown in Figure 2(d). As the a performance metric, we adapt the maximum  
35 number of sound sources that can be rendered in real time.  
36

37 First, we analyze the evolution of the execution time of the CUDA algo-  
38 rithm with the number of sound sources and the effect of parameters such  
39 as the size of RC filters or the number of samples. Figure 5 shows that the  
40 execution time increases linearly with the number of sound sources and that  
41 we can process a large number of sound sources in real time. We use an audio  
42 buffer size of 256 samples, which implies a real-time threshold of 5.8 ms (sam-  
43 pling frequency of 44,1 kHz). The figure shows that the size of the applied RC  
44 filters only affects the execution time slightly and allows us to process around  
45 100 sources in real time using filters with lengths that range from 256 to 1536  
46 coefficients. It is important to highlight that there is an upper bound on the  
47 number of sound sources that can be processed due to the memory size of the  
48 GPU. As illustrated in the figure, the use of RC filters with a length of 1,536  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65



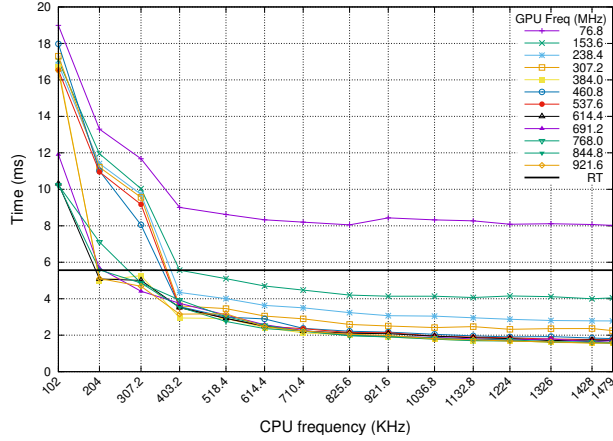
**Fig. 5** Evolution of the execution time of the CUDA implementation with the number of sources using different sizes of the RC filters and an audio buffer of 256 samples at a sample frequency of 44.1 kHz

fills the GPU memory when the implementation handles more than 110 sound sources.


We may not be interested in processing a very large number of sound sources in real time, but instead in processing a smaller number of sound sources in real time while reducing the energy consumption. To this end we can diminish the frequency of the GPU as well as that of the CPU core that is used as host for the CUDA program. To obtain the results in the following figures, we used an audio buffer consisting of 256 samples, RC filters of size 1,024 and 16 loudspeakers. For example, Figure 6 shows that we can process 10 sources in real time even when reducing substantially the GPU and CPU frequencies. Only if we set the lowest frequency of the GPU (76.8 MHz), we cannot process 10 sources in real time for any CPU frequency. However, we can reduce the frequency of the GPU to its second lowest value and process 10 sound sources in real time even with a very small CPU frequency: 403.2 kHz. In general terms, given a fixed GPU frequency, initially the execution time of the algorithm decreases quickly when we increase the CPU frequency from its lowest value. However, when we increase the CPU frequency above 403.2 kHz, the execution time decreases very slowly.

A deeper understanding of the effect of the GPU and CPU frequencies on the computational efficiency of the algorithm can be obtained by analyzing the maximum number of sound sources that can be processed in real time for every GPU-CPU frequencies combination. Figure 7 shows that the maximum number of sound sources increases both with the CPU-GPU frequencies. Once more, it increases very quickly for the lowest frequencies, and more slowly as we approach the highest frequencies.

We are not only interested in maximizing the number of sound sources that can be rendered in real time, but also in reducing the energy consumption of




**Fig. 6** Execution time of the CUDA implementation to process 10 sources with a buffer size of 256 samples and RC filters of size 1024. The horizontal line represents the time to process the sources in real time.



**Fig. 7** Maximum number of sources processed in real time for the different combinations of GPU and CPU frequencies. Values are only shown when at least one source can be processed in real time.

the system. Therefore, we next analyze the energy consumed by the different configurations in order to maximize the sources-per-Watt that can be rendered in real time.

We obtain the energy consumed by the application by performing the product of the power dissipated by the whole platform and the execution time. In all cases, we report the “application energy consumption” obtained by subtracting the observed idle power at minimum CPU-GPU frequencies from the total energy consumption observed during the application execution. We prefer to analyze its energy efficiency by using the sources-per-Watt metric. Figure 8 shows that the maximum number of sources that we can process per Watt increases slowly with the CPU-GPU frequencies. This parameter is between



**Fig. 8** Energy efficiency of the CUDA code computed as the maximum number of sources-per-Watt that can be processed in real time. Values are only shown when at least one source can be processed in real time.

10 and 11 for an ample range of the highest frequencies of both components of the platform, which shows that it is not necessary to use the highest CPU-GPU frequencies to maximize the energy efficiency of the WFS application in the Jetson Nano.

## 5 Conclusions

In this paper we have shown that it is possible to realize a flexible and efficient distributed spatial audio system based on WFS using low-power SoC devices as audio nodes. Specifically, we have leveraged the GPU of several Jetson Nano boards to distribute the audio processing of the sound sources.

We have shown that the application of the NTP synchronization protocol at the acoustic nodes allows us to render spatial sound with sufficient quality, as asserted by the subjective tests and the delay measurements among nodes.

A thorough analysis of the performance of the Jetson Nano device shows that it can process more than 100 of sound sources in real time even using room compensation filters with a huge number of coefficients. We have leveraged the possibility of reducing the frequencies of the GPU and CPU components of the Jetson Nano to reduce the energy consumption of the application. Employing lower frequencies increases very slowly the processing time. The results show that we can reach the maximum number of sources-per-Watt even reducing substantially the frequencies of both components.

## Acknowledgements

This work has been supported by the Spanish Government through TIN2017-82972-R, ESP2015-68245-C4-1-P, the Valencian Regional Government through PROMETEO/2019/109 and the Universitat Jaume I through UJI-B2019-36.

## References

1. A. Berkhout, "A holographic approach to acoustic control," *J. Audio Engineering Society*, vol. 36, pp. 2764–2778, May 1988.
2. D. Theodoropoulos, G. Kuzmanov, and G. Gaydadjiev, "Multi-core platforms for beamforming and Wave Field Synthesis," *IEEE Transactions on Multimedia*, vol. 3, no. 2, pp. 235–245, April 2011.
3. J. A. Belloch, A. Gonzalez, E. S. Quintana-Ortí, M. Ferrer, and V. Välimäki, "GPU-Based Dynamic Wave Field Synthesis Using Fractional Delay Filters and Room Compensation," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 25, no. 2, pp. 435–447, Feb 2017.
4. S. Spors, H. Buchner, and R. Rabenstein, "Efficient active listening room compensation for Wave Field Synthesis," in *Proceedings of the 116th AES Convention*, Berlin, Germany, May 2004.
5. J. Lopez, A. Gonzalez, and L. Fuster, "Room compensation in wave field synthesis by means of multichannel inversion," in *IEEE workshop on Applications of Signal Processing to Audio and Acoustics, 2005.*, oct. 2005, pp. 146 – 149.
6. NVIDIA Corp., "NVIDIA Jetson Linux Developer Guide. PR-06076-R32," May 2020.
7. D. Puccinelli and M. Haenggi, "Wireless sensor networks: applications and challenges of ubiquitous sensing," *Circuits and Systems Magazine, IEEE*, vol. 5, no. 3, pp. 19–31, 2005.
8. W.-P. Chen, J. Hou, and L. Sha, "Dynamic clustering for acoustic target tracking in wireless sensor networks," *Mobile Computing, IEEE Transactions on*, vol. 3, no. 3, pp. 258–271, 2004.
9. A. Berkhout, D. de Vries, and P. Vogel, "Acoustic control by wave field synthesis," *J. Acoustic. Soc. Amer*, vol. 93, pp. 2764–2778, May 1993.
10. S. Spors, A. Kuntz, and R. Rabenstein, "An approach to listening room compensation with wave field synthesis," in *Proc. 24th AES Conference*, Banff, Canada, May 2003.
11. L. Fuster, J. J. Lopez, A. Gonzalez, and P. Faus, "Time and frequency domain room compensation applied to wave field synthesis," in *Proc. Conference on Digital Audio Effects (DAFx-05)*, Madrid, Spain, September 2005.
12. D. Mills, J. Martin, J. Burbank, and W. Kasch, "RFC 5905: Network time protocol version 4: Protocol and algorithms specification," *Internet Engineering Task Force*, 2010.
13. "IEEE standard for a precision clock synchronization protocol for networked measurement and control systems," *IEEE Std 1588-2019 (Revision of IEEE Std 1588-2008)*, pp. 1–499, 2020.
14. T. Neagoe, V. Cristea, and L. Banica, "NTP versus PTP in computer networks clock synchronization," in *2006 IEEE International Symposium on Industrial Electronics*. IEEE, jul 2006.
15. M. A. Lombardi, L. M. Nelson, A. N. Novick, and V. S. Zhang, "Time and frequency measurements using the global positioning system," *Cal Lab: International Journal of Metrology*, vol. 8, no. 3, pp. 26–33, 2001.
16. P. Fubin, Y. Yubo, G. Lei, and S. Liangliang, "The accuracy of IEEE 1588 time synchronization protocol and its improvement," in *2015 12th IEEE International Conference on Electronic Measurement & Instruments (ICEMI)*. IEEE, jul 2015.
17. S. Barrachina, M. Barreda, S. Catalán, M. F. Dolz, G. Fabregat, R. Mayo, and E. Quintana-Ortí, "An integrated framework for power-performance analysis of parallel scientific workloads," *Energy*, pp. 114–119, 2013.