The potential of high-performance computing for the Internet of Sounds

Luca Turchet
Dept. of Information Engineering and Computer Science
University of Trento
Trento, Italy
luca.turchet@unitn.it

Flavio Vella
Dept. of Information Engineering and Computer Science
University of Trento
Trento, Italy
flavio.vella@unitn.it

Sandro Luigi Fiore
Dept. of Information Engineering and Computer Science
University of Trento
Trento, Italy
sandro.fiore@unitn.it

2023 4th International Symposium on the Internet of Sounds | 979-8-3503-8254-9/23/$31.00 ©2023 IEEE | DOI: 10.1109/IEEECONF59510.2023.10335270

***Abstract*—High-Performance Computing (HPC) technology is impacting several industries, including the creative industries and those operating in the Internet of Things. However, thus far, scarce attention has been devoted by researchers to the use of such technology in Internet of Sounds (IoS) settings. The IoS is a new area emerging in industry and academia as an extension of the IoT to the sonic domain, both in musical and non-musical contexts. The IoS relates to a distributed network of sound things, which are devices capable of sensing, acquiring, processing, actuating, and exchanging data serving the purpose of communicating sound-related information. In this position paper, we investigate the integration of HPC technology with the IoS. We illustrate such integration via a set of application scenarios showing how the IoS can benefit from HPC. Finally, we discuss the open challenges ahead of us to implement a successful integration of these two kinds of technology.**

***Index Terms*—Internet of Sounds, High-Performance Computing, audio processing**

I. Introduction

The Internet of Sounds (IoS) is an emerging area of research positioned at the confluence of Internet of Things [1] and Sound and Music Computing [2]. The recent study reported in [3] has proposed a vision for such a field, stating that the IoS relates to the network of Sound Things, i.e., devices capable of sensing, acquiring, processing, actuating, and exchanging data serving the purpose of communicating sound-related information. The field can be seen as the union of the two emerging paradigms of the Internet of Musical Things [4] and the Internet of Audio Things [5], which respectively address musical and non-musical sounds in networked contexts. In the IoS paradigm, heterogeneous devices dedicated to musical and non-musical tasks can interact and cooperate with one another, as well as with other things connected to the Internet, to facilitate sound-based services and applications that are globally available to the users.

Sound Things have the ability to generate large quantities of data, including audio files and associated metadata, features extracted from audio signals, or data about users. For instance, smart musical instruments are envisioned to collect on a daily basis different kinds of data from musicians during their musical activities [6]. The same occurs for wireless acoustic sensor networks (WASNs) used for the monitoring of wildlife [7]–[9] or urban soundscapes in smart cities [10], [11]. Such data requires appropriate computing infrastructures to be processed and analyzed, e.g., for classification or prediction tasks typical of machine learning algorithms or for information inference derived from the application of big data analysis techniques. In a similar vein, some audio processes such as complex physical modelling synthesis or acoustic room simulations are computationally prohibitive to be realized with conventional hardware, especially in real-time and in networked contexts.

High performance computing (HPC) represents an appropriate technology to cope with such tasks, which can typically be solved with parallel computational models (i.e., where different computations can be made concurrently). Indeed, HPC technology has the ability of carrying out large scale computations to solve complex problems. Typically, these problems require to process large amounts of data, and/or need vast availability of computing power to achieve a solution. HPC technology is becoming more and more readily available for academic and industrial applications [12]. Recent years have witnessed a skyrocketing interest towards HPC in many areas, such as healthcare [13], climate studies [14], or space exploration [15], and in general in various applications dealing with complex, parallelizable computations as well as the processing of high quantities of data. Comparatively, however, there has been much less research about the use of HPC in audio contexts, and this is especially true for audio in networked settings. While some scholars have focused on the application of HPC to audio and music tasks (see e.g., [16]–[18]), the HPC literature in the IoS space is rather scarce. Besides the conventional use case of training and optimization of neural models that target audio processing, analysis, and synthesis on remote GPU/TPU clusters, the potentialities of HPC for networked audio applications are largely unexplored. This is particularly true for real-time applications.

To bridge this gap, this paper provides an overview of the opportunities resulting from the integration of HPC technologies in the IoS. We illustrate such integration via a set of application scenarios showing how the IoS can benefit from HPC. Finally, we outline the open research directions in this promising area. With this study, we aim to highlight the potential that HPC offers to the IoS field, but also reflect on the challenges ahead of us.

II. Related Work

A. High-performance computing

HPC is a hardware and software technology that enables running applications on computing systems which feature a large number of interconnected processors. Specifically, these systems employ multiple processors or nodes, interconnected by high-speed networks, to execute tasks concurrently. By dividing the workload among multiple processors, HPC systems can achieve remarkable performance gains and significantly reduce the time required for complex simulations, data analysis, and modeling. The applications of HPC technology span various domains, including applications such as weather forecasting, climate modeling, drug discovery, genomics, aerospace engineering, financial modeling, and many more. HPC systems are particularly useful for problems that involve massive amounts of data, simulation, and machine learning-based models that require a high level of accuracy.

With the continuous advancements in hardware and software, HPC technology continues to evolve rapidly and with significant and constant improvement of computing, network and storage. Modern HPC systems incorporate accelerators like graphics processing units (GPUs) [19] or field-programmable gate arrays (FPGAs) [20] to enhance their processing capabilities by enabling the offload part of the computation. Additionally, developments in cloud computing have enabled the provision of HPC resources on demand, making high-performance computing more accessible and cost-effective for organizations of all sizes [21].

B. HPC for audio technology

Various studies have investigated the integration of HPC with audio technology, especially focusing on GPU and FPGA accelerators (see e.g., [16]–[18]). Some studies have discussed the potential of HPC for audio technology [22], in particular for audio engineering [23]. Moore et al. leveraged an HPC cluster for optimizing Ambisonic surround sound decoders for spatial audio applications (i.e., to obtain an optimal set of Ambisonic decoder parameters) showing the advantage in time consumption when using HPC solutions compared to conventional hardware where processes are run in series [24]. In a similar vein, Sumner et al. used HPC resources to optimize the parameters of neural networks (as well their training) dedicated to the computationally intensive task of emulating human sound localization [25]. Such a study showed the viability of such an approach as well as that HPC enables significant speed-ups and reduces the time-to-solution for sound localization compared to conventional hardware.

Relatedly, some authors have investigated the use of GPUs for audio processing tasks. Hamilton and Webb [26] explored the use of GPUs for room acoustics simulation using a finite difference approximation. Belloch et al. investigated the use of GPUs for performing multi-channel filtering [27], while Tsai et al. used GPUs for achieving spectral model synthesis for real-time sound rendering [28]. Bianchi and colleagues explored the use of GPUs for the typical tasks of Digital Audio Workstations, showing that a GPU-based approach can be a valid alternative to the use of CPUs in the computation of audio effects, such as the rendering of audio tracks after mixing and mastering operations, both in real time and offline [29]. Renney and colleagues investigated the limitations of processing audio in the GPU environment by presenting a performance benchmarking suite [30]

The only study positioned at the intersection of HPC and IoS is reported in [31]. Alsina-Pagès and colleagues proposed a wireless acoustic sensor network for real-time audio event detection based on HPC for remote monitoring of behaviour and surveillance. Specifically, authors developed a distributed system involving nodes equipped with high performance general purpose GPUs.

In industrial contexts, the company HPC Music[1] has developed a dedicated HPC hardware (i.e., a desktop supercomputer), specifically conceived for musical applications. Nevertheless, the primary use of such hardware was devised for co-located computations, not for networked ones.

III. Opportunities for integrating HPC with IoS

The emergence of highly parallel processors and of large-scale multi-core platforms offers the possibility of running networked audio processes previously disregarded as too computationally expensive, in real-time. By leveraging the significant computational power and parallel processing capabilities of HPC technology, the Internet of Sound can achieve breakthroughs in audio processing, analysis, synthesis, and transmission. HPC enables real-time processing, enhances accuracy and efficiency, and opens up unprecedented possibilities for immersive audio experiences and innovative sound-related applications across diverse domains. These include entertainment, communication, healthcare, transportation, and beyond. The following is a non-exhaustive list of opportunities that can be envisioned for HPC in IOS contexts.

- **Real-Time Audio Processing:** HPC systems can handle the computational demands of complex real-time audio processing, enabling advanced audio algorithms to be executed rapidly. This capability is crucial in applications such as immersive audio, virtual reality, and augmented reality, where sound processing needs to be performed in real-time to create an immersive and interactive audio experience. This is especially true for networked settings.
- **Audio Analysis and Pattern Recognition:** HPC can accelerate complex audio analysis tasks such as audio fingerprinting, speech recognition, music classification, and acoustic event detection. By leveraging parallel processing capabilities, HPC systems can process large audio datasets efficiently, enabling faster and more accurate identification and categorization of sounds. This is particularly important for enabling novel cloud-based services for IoS stakeholders. The workload is similar to other tasks that adopt ML techniques such as inference tasks on edge devices based on large or deep pre-trained models.

[1] https://www.hpcmusic.com/

- **Noise Reduction and Audio Enhancement:** HPC technology can aid in real-time noise reduction and audio enhancement algorithms. By employing sophisticated signal processing techniques, from the edge to the cloud and HPC systems, supercomputing facilities can analyze and suppress unwanted noise, enhance speech intelligibility, and improve audio quality in various applications like teleconferencing, voice assistants, and audio recording.
- **Sound Synthesis and Audio Simulation:** HPC systems can enhance sound synthesis techniques and enable realistic audio simulation. By leveraging high computational power, complex sound synthesis algorithms can be executed in real-time, allowing for the creation of immersive virtual environments, realistic sound effects in gaming, and accurate acoustic simulations for architectural or soundscape design.
- **Data Streaming and Network Optimization:** HPC technology can optimize the streaming and transmission of audio data in the IoS. HPC systems can manage large volumes of audio streams, ensure low-latency data transfer, and enhance network performance. This capability is particularly important for applications such as real-time audio streaming, distributed audio processing, and collaborative audio production. Furthermore, HPC may rely on technologies like high performance networks (e.g., InfiniBand Networking [32]) or high-speed storage, wich can benefit IoS applications, including those based on tasks that are not parallelizable.
- **Scalability and Flexibility:** HPC systems offer scalability, allowing the IoS to handle increasing demands for audio processing and analysis. As the number of connected audio devices and sensors grows, HPC can accommodate the processing requirements and efficiently distribute computational tasks across multiple nodes or clusters, ensuring seamless scalability and adaptability.

IV. Envisioned scenarios

The survey of the literature discussed in Section II-B showed how scarce has been the interaction between the field of HPC and that of audio technology, especially in networked settings. HPC opens the door to new applications in the IoS, making problems feasible that are currently infeasible to run on conventional hardware (e.g., desktop computers or standard servers) and allowing computationally expensive algorithms to run in real-time. To show such opportunities, hereinafter we propose a set of scenarios integrating HPC resources in the IoS.

A. Scenario 1: Advanced WASNs

Thanks to HPC technology it is possible to envision new kinds of WASNs which can involve many more nodes than current ones, even in the order of hundreds. This would enable to cover much larger areas than existing deployments, such as that of a whole city (for urban soundscape monitoring) or an extensive natural landscape (for wildlife monitoring, including the understanding of movements of animals, such as birds). Moreover, these WASNs could have a hybrid architecture merging distributed and centralized features. Specifically, each node would comprise not only the conventional microphone-equipped embedded system, but also an accelerator such as a general purpose GPU. This could be used for extracting low-level audio features leveraging highly parallelizable algorithms such as those of frame-based audio analysis techniques (e.g.,, Fast Fourier Transform, Wavelet Transform or Short Time Fourier Transform). The features locally extracted in real-time can then be sent in a streaming fashion (e.g., every 3 seconds) from each node to a remote central server (hosting a HPC cluster) that will be able to process such amount of data nodes in real-time. Such complex processing could be achieved thanks to machine learning algorithms able to classify acoustic events of interest (e.g., gunshots, car accidents, or birds chirping). The optimization of these models could also be conducted on the central server at regular intervals (e.g., once per week). This includes algorithms for hyperparameter optimization, which relates to the process of searching for the hyperparameters of a machine learning model that achieve the best possible performance. Such optimization is computationally intensive since it requires a full model training run for many combinations of different hyperparameter values.

B. Scenario 2: Context-aware musical instruments and cloud-based proactive services

Smart musical instruments [6] are envisioned to generate and leverage large quantities of data to enable a whole new set of proactive services based on the retrieval of the context surrounding the instrument itself (e.g., who is the player, which music is being played, in which musical activity, with whom, where, when, etc.). Such services could be based on a cloud-based recommendation system and find application in a variety of musical activities, including music learning, composition, performance or recreational music making. To effectively infer the context there is the need of processing, especially in real-time, a large amount of data, which is multimodal in nature: this includes audio files and audio streams; logs of interactions with apps, social networks, websites; logs of interactions of the smart instrument with other smart instruments; data from brain-computer interfaces used to classify the mental and emotional state of the player. Equipping musical instruments with such context-aware features entails the use of HPC for running multimodal machine learning algorithms, especially considering that the data streamed to the cloud-server will arrive from a large set of musical instruments, potentially thousands and geographically distributed worldwide.

C. Scenario 3: Large scale synthesis for audience participation

The ability of HPC of performing concurrently many computations enables novel opportunities for composers willing to exploit IoS architectures to explore new artistic forms. These include new kinds of live music concerts where a large audience (e.g., dozen of thousands of people such as those of stadiums) actively participates in the music creation process, in a dialogue with the musicians on the stage. In this scenario, each audience member could reproduce through their smartphone a set of sounds that are computed on and streamed from an HPC cluster. In particular, HPC resources could be dedicated not only to computing such a massive number of audio streams, but also to use specific synthesis techniques unfeasible on the smartphones themselves (e.g., via complex physical models).

D. Scenario 4: Immersive networked music performances

Thanks to the capabilities of HPC it is possible to imagine new kinds of immersive networked music performances where geographically displaced musicians play together over the network. To render credibly the sensation of sharing the same physical space, it is necessary to extend such systems with spatial audio capabilities, which include room acoustics simulation. With conventional computational methods, this is achieved with algorithms that make approximations of the acoustics of a given room (e.g., by means of sampled impulse responses or delay networks). A higher degree of realism can be achieved by using more complex algorithms that generate in real-time reverb by direct modelling of the physical space (over the full frequency range of interest). Furthermore, it is possible to imagine the use of Wave Field Synthesis for achieving the sound spatialization, for instance for audience members attending remotely a concert. This spatial audio technique involves a large number of loudspeakers whose sound can be computed concurrently, exploiting the computational power of HPC hardware.

V. Open challenges

HPC technology may be beneficial to the IoS in terms of computing power, as shown by the scenarios envisioned in Section III. Nevertheless, the integration of the two fields represents a non-trivial task and faces various challenges which are exacerbated at scale by the inherent big data dimension [33]. In order to reach maturity and achieve efficiency, it is necessary to face many challenges that currently prevent the adoption of such integration. These challenges include the following areas: IoS benchmark, specialized hardware/software architectures, compute continuum, sustainability, democratization, standards and scientific community. Hereinafter we briefly report on them.

Challenge 1: IoS Benchmark

To provide a systematic study of the IoS by HPC experts, characterize it in terms of HPC metrics and get a better understanding about the suitability of HPC systems for IoS applications, a specific benchmark could be designed. It would provide guidance about the design of new specialized architectures as well as the assessment of existing ones regarding IoS workloads. It is expected the benchmark to be community-driven, growing over time by means of novel kernels, according to emerging needs and requirements.

Challenge 2: Specialized hardware and software architectures

There is the need for specialized hardware and software architectures specific to the audio domain in networked contexts, especially for real-time scenarios. Recent trends in deep learning imposed hardware accelerators at the edge mainly for the inference phase. These devices require the design of Multi-Chip Modules [34] able to minimize the latency for real-time applications. Although ASIC-based accelerators, Neural Processing Units based CMOS technology are still dominating the market, recent studies are showing promising results in newer and less power-hungry devices. Such technologies including for example neuromorphic device [35] and non-volatile in-memory computing devices based on memristors are at pre-silicon maturity [34]. Investigation of such technologies in the context of IoS may open new scenarios. If the integration of such AI/HPC technology represents a challenge *per se*, still the integration of such computation units with ultra-reliable low-latency wireless networks, in order to minimize the latency of the exchanged sound-based information is even more urgent in the context of IoS. On the software side, the low level of the stack (micro-architectures and compilers) does not require specific studies for IoS workload, however, the optimization of libraries, tools and framework for the inference should consider the specific workload and data characterization. This, therefore, may include the investigation of existing Machine Learning compilers (e.g., microTVM [36]) and auto-adaptive techniques [37] and the integration of other latency in a federated context (e.g., network latency).

Challenge 3: Compute continuum

The wide spectrum, complexity and heterogeneity of advanced IoS applications demand for more articulated architectural solutions and software stacks; in this respect, the Compute Continuum can offer great flexibility to IoS developers and final users, by relying on hybrid IoT-edge-cloud infrastructures. In such a multifaceted, continuum ecosystem, the HPC can play a key role as enabler of larger-scale and real-time IoS scenarios, while bringing at the same time additional complexity and technological challenges to guarantee the proper level of transparency and abstraction. However here, we need to face important challenges. First how to integrate emerging HPC and computing paradigms for memory-centric applications into IoS workloads such as Processing in-memory architecture [38] and in-network computing [39] may also have a disruptive impact on IoS. Both these newer computing paradigms avoid moving the computation inside the processors following the traditional Von Neumann model. Processing in-memory enables simple computation inside the memory (SRAM or DRAM) by keeping the data where it resides. In-network compute offloads small computations into network devices such as Smart Network Interface Cards. Both these two methodologies proved to be effective in applications where the role of latency plays a fundamental role in the performance. Secondly, the IoS' highly distributed architecture poses many challenges in the definition of a suitable automatic mechanism for data processing [40]. These challenges include for example the definition of data analytics pipelines, their orchestration and optimization with respect to audio-based tasks, which account for the distributed nature of the nodes of IoS applications, the integration of the data and the pre-processing.

Challenge 4: Sustainability

A very important challenge underlying the HPC field is the power consumption and the related environmental impact, and its application in IoS contexts is no exception. There is the need for optimizing the computing resources so to achieve a high level of energy efficiency. Various researchers in the HPC field have proposed methods for improving power management schemes and thus the resulting carbon footprint [41], [42]. Such methods can be adapted to the specific IoS scenarios towards the achievement of a more sustainable IoS recently advocated by some researchers [43]. In this context, the adoption of novel computing architectures and models, as previously described, yields significant benefits in terms of power consumption. These advancements result in notable gains in energy efficiency, mitigating power-related concerns. By leveraging cutting-edge hardware and new materials for semiconductor fabrication, embracing new parallel processing paradigms, employing low-power components, and implementing effective energy management techniques, modern computing systems are demonstrating remarkable reductions in power usage [34], [44].

Challenge 5: Democratization

Unfortunately, getting access to an HPC facility can represent a barrier or can be limited to a very small group of IoS scientists. In this respect, digital infrastructure initiatives, i.e., EOSC (European Open Science Cloud) [45], could help providing the needed HPC capacity to overcome this issue and open (democratize) the access to the wider IoS research community as a whole.

Challenge 6: Lack of standards

Solutions in HPC for audio in networked settings are yet to be extensively developed. To avoid fragmentation of this emerging area, proper standards could be defined, which may significantly contribute to the adoption of such solutions. Thus the attention of industry and academia could be devoted to define such standards (e.g., interoperable protocols), in order to optimally govern the interaction between Sound Things, HPC, and IoS stakeholders.

Challenge 7: Scientific Community

The IoS research field has a great potential in terms of expansion over the next decade. To move faster and grow consistently, community building will be key both in terms of collaborations and networking, but also in terms of link with large digital infrastructures. In this respect and according to the EU Data Strategy [46], the development of a "Data space for the IoS" will provide a unique opportunity in terms of digital ecosystem for IoS scientists. Such a data space will (i) build around the data collected in this field, fostering sharing and re-use of data according to Open Science and FAIR [47] principles, and (ii) support large-scale data analytics by leveraging the HPC/cloud compute capacity provided by relevant e-infrastructural initiatives (i.e., EOSC [45], PRACE [48]).

VI. Concluding Remarks

This paper proposed a synthesis of two different technologies, HPC and IoS. This synthesis was illustrated via a number of application scenarios that could be implemented in the short- and medium-term. Of course, there are other aspects that may have a significant impact on the intersection between IoS and HPC field. This includes, for example, the role of the IoS network in terms of Wireless Sensor Network and its properties among others. The proposed integration of HPC in IoS contexts represents a rather novel domain of study that explores the potential for new advanced processor architectures to transform the current landscape of audio in networked contexts.

The IoS can benefit from the functionalities provided by HPC, which may help to further develop and improve current IoS applications. Even though the convergence of HPC and IoS brings several opportunities to improve the audio and musical industry, there are many research challenges that must be addressed before the potential of this integration can be fully unleashed. Such challenges include technical (e.g., latency) and sustainability (e.g., energy efficiency) aspects.

With this position paper, our aim was primarily to raise attention towards the yet unexplored potential of using HPC resources in IoS contexts. This research topic is still in a preliminary stage and calls for more research capable of providing effective solutions to bridge the current gaps. It is hoped that the current study could spur the interest of researchers in implementing the envisioned integration, thus advancing the state of the art of the IoS field.

References

[1] L. Atzori, A. Iera, and G. Morabito, "The internet of things: a survey," *Computer networks*, vol. 54, no. 15, pp. 2787–2805, 2010.

[2] G. Widmer, D. Rocchesso, V. Välimäki, C. Erkut, F. Gouyon, D. Pressnitzer, H. Penttinen, P. Polotti, and G. Volpe, "Sound and music computing: Research trends and some key issues," *Journal of New Music Research*, vol. 36, no. 3, pp. 169–184, 2007.

[3] L. Turchet, M. Lagrange, R. C., G. Fazekas, N. Peters, J. Østergaard, F. Font, T. Bäckström, and C. Fischione, "The internet of sounds: Convergent trends, insights and future directions," *IEEE Internet of Things Journal*, 2023.

[4] L. Turchet, C. Fischione, G. Essl, D. Keller, and M. Barthet, “Internet of Musical Things: Vision and Challenges,” *IEEE Access*, vol. 6, pp. 61 994–62 017, 2018.

[5] L. Turchet, G. Fazekas, M. Lagrange, H. Shokri Ghadikolaei, and C. Fischione, “The internet of audio things: state-of-the-art, vision, and challenges,” *IEEE Internet of Things Journal*, vol. 7, no. 10, pp. 10 233–10 249, 2020.

[6] L. Turchet, “Smart Musical Instruments: vision, design principles, and future directions,” *IEEE Access*, vol. 7, pp. 8944–8963, 2019.

[7] S. S. Sethi, R. M. Ewers, N. S. Jones, C. D. L. Orme, and L. Picinali, “Robust, real-time and autonomous monitoring of ecosystems with an open, low-cost, networked device,” *Methods in Ecology and Evolution*, vol. 9, no. 12, pp. 2383–2387, 2018.

[8] I. Zualkernan, J. Judas, T. Mahbub, A. Bhagwagar, and P. Chand, “An aiot system for bat species classification,” in *2020 IEEE International Conference on Internet of Things and Intelligence System (IoTaIS)*. IEEE, 2021, pp. 155–160.

[9] E. Rovithis, N. Moustakas, K. Vogklis, K. Drossos, and A. Floros, “Design recommendations for a collaborative game of bird call recognition based on internet of sound practices,” *Journal of the Audio Engineering Society*, vol. 69, no. 12, pp. 956–966, 2021.

[10] J. Segura-Garcia, S. Felici-Castell, J. J. Perez-Solano, M. Cobos, and J. M. Navarro, “Low-cost alternatives for urban noise nuisance monitoring using wireless sensor networks,” *IEEE Sensors Journal*, vol. 15, no. 2, pp. 836–844, 2014.

[11] J. Ardouin, L. Charpentier, M. Lagrange, F. Gontier, N. Fortin, D. Ecotière, J. Picaut, and C. Mietlicky, “An innovative low cost sensor for urban sound monitoring,” in *INTER-NOISE and NOISE-CON Congress and Conference Proceedings*, vol. 258, no. 5. Institute of Noise Control Engineering, 2018, pp. 2226–2237.

[12] M. Golasowski, J. Martinovič, M. Levrier, S. Hachinger, S. Karagiorgou, A. Papapostolou, S. Mouzakitis, I. Tsapelas, M. Caballero, M. Aldinucci *et al.*, “Toward the convergence of high-performance computing, cloud, and big data domains,” in *HPC, Big Data, and AI Convergence Towards Exascale*. CRC Press, 2022, pp. 1–16.

[13] D. Oniga, B. Cantalupo, E. Tartaglione, D. Perlo, M. Grangetto, M. Aldinucci, F. Bolelli, F. Pollastri, M. Cancilla, L. Canalini *et al.*, “Applications of ai and hpc in the health domain,” in *HPC, Big Data, and AI Convergence Towards Exascale*. CRC Press, 2022, pp. 217–239.

[14] S. Fiore, D. Elia, C. Palazzo, F. Antonio, A. D’Anca, I. Foster, and G. Aloisio, “Towards high performance data analytics for climate change,” in *High Performance Computing: ISC High Performance 2019 International Workshops*. Springer, 2019, pp. 240–257.

[15] S. Esposito, C. Albanese, M. Alderighi, F. Casini, L. Giganti, M. L. Esposti, C. Monteleone, and M. Violante, “Cots-based high-performance computing for space applications,” *IEEE Transactions on Nuclear Science*, vol. 62, no. 6, pp. 2687–2694, 2015.

[16] T. Risset, R. Michon, Y. Orlarey, S. Letz, G. Müller, and A. Gbadamosi, “faust2fpga for ultra-low audio latency: Preliminary work in the syfala project,” in *IFC 2020-Second International Faust Conference*, 2020, pp. 1–9.

[17] C. Wegener, S. Stang, and M. Neupert, “FPGA-accelerated real-time audio in pure data,” in *Proc. Int. Conf. in Sound and Music Computing, SMC-22*, 2022.

[18] M. Popoff, R. Michon, T. Risset, Y. Orlarey, and S. Letz, “Towards an fpga-based compilation flow for ultra-low latency audio signal processing,” in *SMC-22-Sound and Music Computing*, 2022.

[19] W. J. Dally and J. Nickolls, “The gpu computing era,” *IEEE Micro*, vol. 30, no. 02, pp. 56–69, mar 2010.

[20] T. Nguyen, C. MacLean, M. Siracusa, D. Doerfler, N. J. Wright, and S. Williams, “Fpga-based hpc accelerators: An evaluation on performance and energy efficiency,” *Concurrency and Computation: Practice and Experience*, vol. 34, no. 20, p. e6570, 2022. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1002/cpe.6570

[21] M. A. Netto, R. N. Calheiros, E. R. Rodrigues, R. L. Cunha, and R. Buyya, “Hpc cloud for scientific and business applications: taxonomy, vision, and research challenges,” *ACM Computing Surveys*, vol. 51, no. 1, pp. 1–29, 2018.

[22] R. Dobson, R. Bradford *et al.*, “High performance audio computing: a position paper,” in *Proceedings of the International Computer Music Conference*, 2008, pp. 213–216.

[23] D. J. Moore and J. P. Wakefield, “The potential of high performance computing in audio engineering,” in *Proceedings of the 126th Audio Engineering Society Convention*, 2009.

[24] D. Moore, Q. Mair, and J. P. Wakefield, “Ambisonic audio system optimization using a hpc cluster,” in *Proceedings of the 9th International Conference on Manufacturing Research*, 2011.

[25] E. M. Sumner, M. Aach, A. Lintermann, R. Unnthorsson, and M. Riedel, “Speed-up of machine learning for sound localization via high-performance computing,” in *Proceedings of the 26th International Conference on Information Technology*. IEEE, 2022, pp. 1–4.

[26] B. Hamilton and C. J. Webb, “Room acoustics modelling using gpu-accelerated finite difference and finite volume methods on a face-centered cubic grid,” *Proc. Digital Audio Effects (DAFx), Maynooth, Ireland*, pp. 336–343, 2013.

[27] J. A. Belloch, B. Bank, L. Savioja, A. Gonzalez, and V. Välimäki, “Multi-channel IIR filtering of audio signals using a GPU,” in *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2014, pp. 6692–6696.

[28] P.-Y. Tsai, T.-M. Wang, and A. Su, “GPU-based spectral model synthesis for real-time sound rendering,” in *Proceedings of the 13th International Conference on Digital Audio Effects*, 2010, pp. 1–5.

[29] D. Bianchi, F. Avanzini, A. Baratè, L. A. Ludovico, and G. Presti, “A GPU-Oriented Application Programming Interface for Digital Audio Workstations,” *IEEE Transactions on Parallel and Distributed Systems*, vol. 33, no. 8, pp. 1924–1938, 2021.

[30] H. Renney, B. R. Gaster, and T. J. Mitchell, “There and back again: The practicality of GPU accelerated digital audio,” in *Proceedings of the International Conference on New Interfaces for Musical Expression*, 2020.

[31] R. M. Alsina-Pagès, J. Navarro, F. Alías, and M. Hervás, “homesound: Real-time audio event detection based on high performance computing for behaviour and surveillance remote monitoring,” *Sensors*, vol. 17, no. 4, p. 854, 2017.

[32] S. Sur, M. J. Koop, L. Chai, and D. K. Panda, “Performance analysis and evaluation of Mellanox ConnectX InfiniBand architecture with multi-core platforms,” in *15th Annual IEEE Symposium on High-Performance Interconnects*. IEEE, 2007, pp. 125–134.

[33] J. Ejarque, R. M. Badia, L. Albertin, G. Aloisio, E. Baglione, Y. Becerra, S. Boschert, J. R. Berlin, A. D’Anca, D. Elia, F. Exertier, S. Fiore, J. Flich, A. Folch, S. J. Gibbons, N. Koldunov, F. Lordan, S. Lorito, F. Løvholt, J. Macías, F. Marozzo, A. Michelini, M. Monterrubio-Velasco, M. Pienkowska, J. de la Puente, A. Queralt, E. S. Quintana-Ortí, J. E. Rodríguez, F. Romano, R. Rossi, J. Rybicki, M. Kupczyk, J. Selva, D. Talia, R. Tonini, P. Trunfio, and M. Volpe, “Enabling dynamic and intelligent workflows for hpc, data analytics, and ai convergence,” *Future Generation Computer Systems*, vol. 134, pp. 414–429, 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0167739X22001364

[34] C. Silvano, D. Ielmini, F. Ferrandi, L. Fiorin, S. Curzel, L. Benini, F. Conti, A. Garofalo, C. Zambelli, E. Calore *et al.*, “A survey on deep learning hardware accelerators for heterogeneous hpc platforms,” *arXiv preprint arXiv:2306.15552*, 2023.

[35] D. Ielmini and S. Ambrogio, “Emerging neuromorphic devices,” *Nanotechnology*, vol. 31, no. 9, p. 092001, 2019.

[36] C. Liu, M. Jobst, L. Guo, X. Shi, J. Partzsch, and C. Mayr, “Deploying machine learning models to ahead-of-time runtime on edge using microtvm,” *arXiv preprint arXiv:2304.04842*, 2023.

[37] P. S. Labini, M. Cianfriglia, D. Perri, O. Gervasi, G. Fursin, A. Lokhmotov, C. Nugteren, B. Carpentieri, F. Zollo, and F. Vella, “On the anatomy of predictive models for accelerating gpu convolution kernels and beyond,” *ACM Trans. Archit. Code Optim.*, vol. 18, no. 1, jan 2021. [Online]. Available: https://doi.org/10.1145/3434402

[38] H. A. D. Nguyen, J. Yu, M. A. Lebdeh, M. Taouil, S. Hamdioui, and F. Catthoor, “A classification of memory-centric computing,” *J. Emerg. Technol. Comput. Syst.*, vol. 16, no. 2, jan 2020. [Online]. Available: https://doi.org/10.1145/3365837

[39] S. Kianpisheh and T. Taleb, “A survey on in-network computing: Programmable data plane and technology specific applications,” *IEEE Communications Surveys & Tutorials*, 2022.

[40] M. Asch, T. Moore, R. Badia, M. Beck, P. Beckman, T. Bidot, F. Bodin, F. Cappello, A. Choudhary, B. de Supinski, E. Deelman, J. Dongarra, A. Dubey, G. Fox, H. Fu, S. Girona, W. Gropp, M. Heroux, Y. Ishikawa, K. Keahey, D. Keyes, W. Kramer, J.-F. Lavignon, Y. Lu, S. Matsuoka, B. Mohr, D. Reed, S. Requena, J. Saltz, T. Schulthess, R. Stevens, M. Swany, A. Szalay, W. Tang, G. Varoquaux, J.-P. Vilotte, R. Wisniewski, Z. Xu, and I. Zacharov, “Big data and extreme-scale computing: Pathways to convergence-toward a shaping strategy for a future software

and data ecosystem for scientific inquiry," *The International Journal of High Performance Computing Applications*, vol. 32, no. 4, pp. 435–479, 2018.

[41] M. E. Haque, K. Le, Í. Goiri, R. Bianchini, and T. D. Nguyen, "Providing green slas in high performance computing clouds," in *2013 International Green Computing Conference Proceedings*. IEEE, 2013, pp. 1–11.

[42] J. Haj-Yahya, A. Mendelson, Y. B. Asher, and A. Chattopadhyay, *Energy Efficient High Performance Processors: Recent Approaches for Designing Green High Performance Computing*. Springer, 2018.

[43] L. Gabrielli and L. Turchet, "Towards a sustainable Internet of Sounds," in *Proceedings of the International Audio Mostly Conference*, 2022, pp. 231–238.

[44] A. Sachan, P. Srivastav, and B. Ghoshal, "Learning based application driven energy aware compilation for gpu," *Microprocessors and Microsystems*, vol. 94, p. 104664, 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S014193312200196X

[45] D.-G. for Research and E. E. B. Innovation (European Commission), "Strategic Research and Innovation Agenda (SRIA) of the European Open Science Cloud (EOSC)," 2022, https://op.europa.eu/en/publication-detail/-/publication/f9b12d1d-74ea-11ec-9136-01aa75ed71a1/language-en/format-PDF/source-289813286 [Accessed: (21/07/2023)].

[46] EC, "Communication: A european strategy for data," 2020, https://commission.europa.eu/document/b456e96e-5810-4bb3-94ca-16c4d8d911b2_en [Accessed: (21/07/2023)].

[47] M. D. e. a. Wilkinson, "The FAIR Guiding Principles for scientific data management and stewardship," *Scientific Data*, no. 3, 2016.

[48] "PRACE- Partnership for Advanced Computing in Europe," https://prace-ri.eu/ [Accessed: (21/07/2023)].