



HAL
open science

Beam-tracing domain decomposition method for urban acoustic pollution

G. Gbikpi-Benissan, F. Magoules

► **To cite this version:**

G. Gbikpi-Benissan, F. Magoules. Beam-tracing domain decomposition method for urban acoustic pollution. 13th International Symposium on Distributed Computing and Applications to Business, Engineering and Science (DCABES), Nov 2014, Xianning, China. 10.1109/dcabes.2014.34 . hal-01273950

HAL Id: hal-01273950

<https://hal.science/hal-01273950>

Submitted on 20 Oct 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Beam-tracing domain decomposition method for urban acoustic pollution

Guillaume Gbikpi-Benissan*

Frédéric Magoulès*

Abstract

This paper covers the fast solution of large acoustic problems on low-resources parallel platforms. A domain decomposition method is coupled with a dynamic load balancing scheme to efficiently accelerate a geometrical acoustic method. The geometrical method studied implements a beam-tracing method where intersections are handled as in a ray-tracing method. Beyond the distribution of the global processing upon multiple sub-domains, a second parallelization level is operated by means of multi-threading and shared memory mechanisms.

Numerical experiments show that this method allows to handle large scale open domains for parallel computing purposes on few machines. Urban acoustic pollution arising from car traffic was simulated on a large model of the Shinjuku district of Tokyo, Japan. The good speed-up results illustrate the performance of this new domain decomposition method.

Keywords: Domain decomposition methods; Parallel and distributed computing; Acoustics; Ray-tracing methods; Beam-tracing methods

1 Introduction

Important material resources is often a requirement, either for the processing of large volumes of data, or for the fast resolution of large problems. Especially in parallel computing, the effectiveness of the designed algorithms is related to their scaling behavior on massively parallel platforms. However, for a great part of the scientific community, there remains a need for big simulations with very low resources. Then, it could be of major interest to design algorithms with good acceleration properties on very small clusters of machines.

In this study, we tackle the simulation of car traffic noise level at a whole district scale. In a lot of countries, environmental noise has become a major source of stress and of various diseases like hypertension or heart ischaemia. Fast noise simulation could be an effective solution to assess the effect of architectural configurations, in order to reduce the acoustic pollution or to maintain a maximum noise level. Given that sound propagates far away in open domains, a large scale model is required, and well designed methods are a key point to obtain fast simulations.

The more physically correct algorithms are based on numerical methods, such as the Boundary Element Method [43], the Infinite Element Method [4, 3], the Finite Element Method [18, 40], the Stabilized Finite Element [16] and the coupling methods [2]. Even if these methods accurately approximate the mathematical equations of the acoustic problems, they unfortunately need a lot of computational power and memory. Hence, they are not good candidates for our context of low resources. Furthermore, relatively to the wavelength, these methods have some difficulties to handle extremely large domains and open domains, such as ones involved in exterior acoustic problems.

Contrariwise, geometrical methods, such as the image-source method [1, 37], the ray-tracing method [9] and the beam-tracing method [10, 19], are well suited for large domains

*Ecole Centrale Paris, France (correspondence, frederic.magoules@hotmail.com).

and open domains. But only simulations of high frequencies, relatively to the size of the problem, provide accurate results, due to the fact that, in such methods, the sound propagates in straight lines. Fortunately, this is not a limitation for our car traffic noise simulation.

In this study, we consider geometrical methods for the fast simulation of urban acoustic pollution within a large open area. After a brief overview of geometrical methods, we describe a hybrid method coupling beam-tracing and ray-tracing methods, with some similarities to frustum-tracing [6]. Then we present the parallelization of the algorithm by means of techniques from domain decomposition methods [30], widely used for numerical methods in acoustics [28]. Finally we illustrate the efficiency of our parallelization technique, using a few machines for simulations run on a model of the shinjuku district of Tokyo, Japan.

2 Hybrid geometrical acoustic method

Commonly, geometrical methods compute multiple paths followed by the sound emitted from a source, and measure the acoustic pressure at some points called microphones. The image-source method [1, 37] creates virtual sources each time a sound ray reflects on a surface inside the 3D model. Sound pressure level is evaluated at a microphone by adding the contribution of the sources and virtual sources for which there is no obstacle on the direct path toward that microphone.

The ray-tracing method [9] divides the energy of each source between a huge number of elementary particles (similar to their physical counterpart, the photons) which propagate as sound rays. The energy of a particle gradually decline due to air damping, and it also loses energy when it reflects on a surface. Once this energy reaches a threshold low value, the particle is deleted.

The beam-tracing method is based on a similar principle, except that sound rays are replaced by sound beams, then the source energy is split between beams, according to its power in each direction. This method is harder to implement, mainly because intersections between beams and surfaces are more complex to detect and to handle. For example, each time a beam partially hits one single surface, it needs to be split into sub-beams, what can dramatically increase the complexity of the computation, due to the recursion of this principle. However, for the same simulation quality, there is much less beams to launch, relatively to ray-tracing.

Ray-tracing methods have been extensively studied in the context of computer generated imagery. Efficient generation of hierarchical partitioning of good quality [39, 41] for the handling of intersection detection, the use of vectorial instructions [42, 41], and graphics processing unit (GPU) [15, 17] allowed to accelerate the ray-tracing. Large scale models got some attention too [42, 5, 20]. However, these optimized methods are very efficient when the rays follow similar trajectories, which is rarely the case in acoustic analysis. On the other hand, even if beam-tracing generates great quality soft shadows and anti-aliasing in computer graphics [36], the additional complexity and slowdown do not worth it.

The hybrid approach considered here applies beam-tracing principles, except that the intersections between beams and objects surface are handled as in ray-tracing method. In other words, instead of testing the whole volume of the beam against the model, only the guiding ray of the beam is tested. Hence, in this method, the beams are never split on reflection. This trade-off between precision and speed allows to use very efficient ray shooting methods. In practice, the drawback due to the approximation error can sufficiently be mitigated by processing very small beams. Moreover, a beam becomes larger far from the energy source, where its energy is significantly depleted, what would reduce the impact of an approximation error.

3 Parallel computing

Advantage of domain decomposition In a standard implementation, the input and the output of parallel computing are sequential. According to the Amdahl's law, the global speed-up of a parallel program containing a percentage f of sequential part is bounded by $1/(1 - f)$. One of the main goals of the Domain Decomposition Methods (DDM) is to allow the whole program to be parallel. The idea is to split the model into sub-domains, such that the loading, output and result gathering could be done in parallel for each sub-domain. The splitting of the model can be done only once for multiple simulations.

Another important advantage of the DDM is to allow to consider very large models and a great number of microphones. This is particularly important when generating volumetric noise map, since the number of microphones quickly becomes very large. Each process only needs to allocate memory for the sub-domains it is working on instead of allocating memory for the whole domain.

Microphones partitioning In this kind of decomposition, only the set of microphones is split, while the complete geometry is replicated. Each sub-domain is totally independent so the program can be run in parallel without synchronization. Even if a whole part of the work is replicated, especially the intersections detection and the loading/preprocessing of the model, such a decomposition can be efficient for a simple geometry model populated with a huge set of microphones. The absence of communication between the processes is a major advantage in the context of a largely distributed or loosely connected platform. However, there is no easy way to mitigate load balancing issues when the processing time of sub-domains varies a lot.

On demand geometry and microphones loading With prior computation of some hierarchical acceleration structure, sub-domains can be loaded only when needed, but the actual parallelization is done on the set of beams. At the beginning, first level of the hierarchy is loaded, and throughout the computation, this structure is accessed several times. If during the traversal of the structure, a node not yet loaded is reached, the corresponding data are loaded. If it is an interior node, these data are the child nodes information, if it is a leaf, the data are the corresponding mesh and microphone data. Beside the specific precomputed acceleration structure, this model of parallelization is quite hard to implement. The latency for loading a node of the hierarchical structure should be hidden, for instance by changing the current beam to one whose data are ready. The gathering of the final results is also complicated since the data are spread over all processes, and potential copies need to be merged. The main drawback is the risk for each process to load the complete model when beams largely spread inside the area. That is particularly the case in exterior acoustics problems where beams often go far and are widely dispersed.

Geometry and microphones partitioning The original method used in this work matches more closely a domain decomposition approach [30]. Here, both the geometry and the microphones are split into multiple sub-domains. When a beam goes out of a sub-domain, it is sent to the process working on the sub-domain intersected. Interface conditions [22] are used to assure the continuity of the beam properties from one sub-domain to another one. Efficient interface conditions can be designed by a continuous approach [8, 14, 7]. Continuous optimized approaches improve the convergence of the algorithm [13, 11, 24, 23, 21, 11]. Similarly, the performance of the algorithm can be increased by using a discrete optimized approach [31, 35, 34, 33, 32, 38]. In [12], a link is established between these two approaches. In this paper we consider a domain decomposition approach as described in [27, 26, 29] but extended to beam-tracing [25]. In contrast, unlike classical domain decomposition methods, load-balancing issues can not be fixed in a static way, where a one-to-one correspondence is done between the processes and the sub-domains. Indeed, considering for instance the case where there would be only one sound source, the sub-domain containing the source would

have the most computational load. This is why a more complex load balancing scheme has to be used. The idea is that each process starts with one sub-domain, but when it has few remaining beams to handle, it starts loading one or more sub-domains still containing a lot of unhandled beams. The number of sub-domains which can be loaded simultaneously is limited by the memory size. In order to overlap the results gathering time, sub-domains can also be unloaded when few beams remain to be processed. Since the gathering operation mainly uses the communication system, this overlapping does not slow down the computation. Nonetheless, the loading and unloading of sub-domains take a long time, hence it would be preferable to avoid such operations as much as possible. For that, it is still relevant to approximate at best an average good load balancing. The solution used is to process first the exchanged beams and then to equalize the number of remaining beams in each sub-domain comparatively to the processing power affected to this sub-domain by loading or unloading sub-domains only when this ratio deviates too much from the average.

At last, the processing of a sub-domain itself is parallelized by coupling a shared memory multi-threading and a work-sharing load balancing. This allows an additional acceleration of the computation without any data replication. There is no need of gathering local outputs at the end since the output data are shared. However, it required to implement a thread-safe access mechanism. A simple fine-grained locking were implemented by setting a spin-lock mutex on each cell of the output array. Given that the multi-threaded beam-tracing generates few contention, the busy-waiting time is insignificant, hence the spin-lock is an acceptable trade-off between simplicity and efficiency for our acoustic problem.

4 Results and discussions

In this study, we tackled the simulation of urban acoustic pollution related to car traffic. Experiments were run on a model of the Shinjuku district of Tokyo, Japan. It is an area of 2.5 km by 1.5 km around the point of coordinates $35^{\circ}41'23''N$, $139^{\circ}1'25''E$. By hosting the busiest train station of the world, multiple commercial and administrative buildings including the administration center of the government of Tokyo, and around ten skyscrapers at least 200 meters high, the Shinjuku district faces a large number of activities generating a very important car traffic. Figure 1 shows a global view of the simulation model. The whole district buildings has been represented, and as we can see on figures 2 and 3, it is a detailed designing leading to a quite large model.

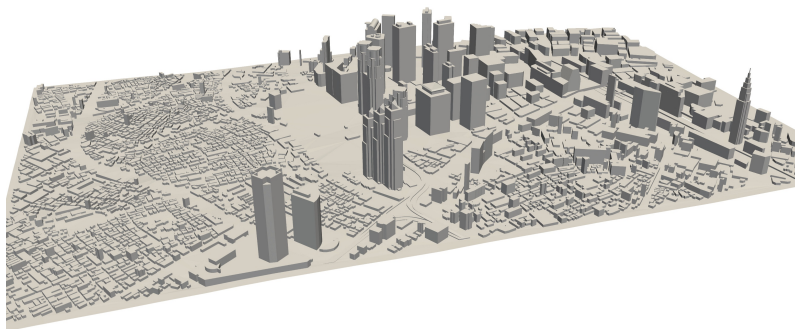


Figure 1: Global view of the virtual model of Shinjuku, district of Tokyo, Japan.

As we were saying at the beginning, our goal is to study the efficiency of the proposed DDM-based hybrid geometrical method on low-computer resources architectures. The numerical experiments were run on a hybrid, both distributed and shared memory, computa-




Figure 2: Close view of the virtual model of Shinjuku, just next to the Tokyo Metropolitan Government Building.




Figure 3: Close view of the virtual model of Shinjuku, under the Metropolitan Expressway No. 4.

tional platform consisting of 8 workstations, each one containing two quad core processors (a total of 64 cores). Each machine was provided with 8 Gigabytes RAM (Random Access Memory) and a Windows 7 operating system. The volume of the hexahedral bounding box of the model was populated with a total number of 13.7 millions of microphones, regularly distributed on a 3D grid. According to the dimensions of the district, that corresponds to a microphone every 4 meters in each of the three spatial directions. We simulated 49 sound sources, most of them placed at major crossroads, and each one launching 3 millions of beams. Figure 4 illustrates the noise level distribution obtained from such a simulation.

Speed-ups are computed using the total execution time, including the loading of the model, the acceleration structure generation and the saving of the results. A speed-up S is determined by $S = t_{ref}/t$, where t_{ref} is the execution time of the sequential version, and t , the time consumed by the parallel version. Table 1 shows different speed-ups relative to various DDM/multi-threading schemes. The computation was distributed such that each core runs exactly 2 threads, which corresponds to 16 threads per workstation. Then, for the 16, 32, 64 and 128 threads, we used respectively 1, 2, 4 and 8 workstations. Even if we are not in a scaling study, one could remark that the classical parallelization without DDM underperforms from 64 threads (4 nodes) to 128 threads (8 nodes). The use of multiple

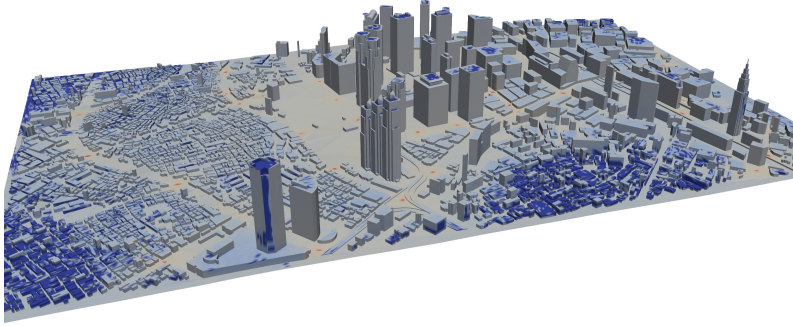


Figure 4: Simulation of a sound pressure level distribution in Shinjuku, district of Tokyo, Japan.

Table 1: Speed-up of the acoustic DDM program (Ethernet) with respect to the number of threads and sub-domains.

	16 threads	32 threads	64 threads	128 threads
1 sub-domain	16.03	25.20	30.26	23.19
2 sub-domains	16.28	30.64	49.25	56.64
4 sub-domains	16.15	29.36	50.82	54.26
8 sub-domains	16.24	29.16	52.79	83.72

sub-domains allows a bigger part of the application to be parallelized, and gives better results. Let us consider a sequential version consuming about 4 hours to complete the whole simulation. With only 4 workstations, the simple multi-threaded simulation would last about 8 minutes (speed-up for 1 sub-domain and 64 threads), while the version with 8 sub-domains would not exceed 5 minutes. However, that is the limit for the classical hybrid distributed/shared-memory parallelization. According to the last result in the table, the DDM approach would allow to reduce the maximum execution time from 4 hours to 3 minutes, using 8 workstations composed of 2 quad-core processors with 8 GB RAM. Table 2 shows parallel execution times (in seconds) for our simulation where the sequential processing lasted 5392 seconds (1 hour and 30 minutes).

Table 2: Execution time (in seconds) of the acoustic DDM program (Ethernet) with respect to the number of threads and sub-domains.

	16 threads	32 threads	64 threads	128 threads
1 sub-domain	336,31	213,96	178,18	232,53
2 sub-domains	331,24	175,98	109,47	95,20
4 sub-domains	333,97	183,68	106,09	99,37
8 sub-domains	332,11	184,93	102,14	64,40

5 Conclusion

Acoustic simulation software is an invaluable tool to handle noise problems, both inside closed and open areas. In this paper, we reminded some principal geometrical methods which basically consider the sound propagation as straight line paths shot from a sound source.

An hybrid beam/ray-tracing algorithm was parallelized by means of domain decomposition methods coupled with a dynamic load balancing. An additional acceleration was gained by a second level parallelization based on multi-threading in shared memory environment.

Numerical simulations have been carried out to evaluate the performance and efficiency of this new method on low-resources platforms, for analyzing large scale urban acoustic pollution. This new method allows to obtain a significant acceleration on few processors by parallelizing the whole algorithm, including the gathering process. As our experiments show, it can be used on inexpensive systems, such as a cluster of workstations provided with an Ethernet interconnection, since it reduces the memory usage and the communication bandwidth.

References

- [1] J. B. Allen and D. A. Berkley. Image method for efficiently simulating small-room acoustics. *The Journal of the Acoustical Society of America*, 65(4):943–950, 1979.
- [2] J.-C. Autrique and F. Magoulès. Numerical analysis of a coupled finite-infinite element method for exterior Helmholtz problems. *Journal of Computational Acoustics*, 14(1):21–43, 2006.
- [3] J.-C. Autrique and F. Magoulès. Studies of an infinite element method for acoustical radiation. *Applied Mathematical Modelling*, 30(7):641–655, 2006.
- [4] J.-C. Autrique and F. Magoulès. Analysis of a conjugated infinite element method for acoustic scattering. *Computers and Structures*, 85(9):518–525, 2007.
- [5] G. Cadet, S. Zambal, and B. Lécussan. Rendering complex scenes on clusters with limited precomputation. In *Proceedings of the International Symposium on High Performance Computational Science and Engineering*, volume 172, pages 79–96. Springer Boston, 2005.
- [6] A. Chandak, C. Lauterbach, M. Taylor, Z. Ren, and D. Manocha. AD-Frustum: Adaptive Frustum Tracing for Interactive Sound Propagation. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1707–1722, 2008.
- [7] P. Chevalier and F. Nataf. Symmetrized method with optimized second-order conditions for the Helmholtz equation. In *Domain decomposition methods, 10 (Boulder, CO, 1997)*, pages 400–407. Amer. Math. Soc., Providence, RI, 1998.
- [8] B. Després. Domain decomposition method and the Helmholtz problem.II. In *Second International Conference on Mathematical and Numerical Aspects of Wave Propagation (Newark, DE, 1993)*, pages 197–206, Philadelphia, PA, 1993. SIAM.
- [9] D. O. Elorza. Room acoustics modeling using the raytracing method: implementation and evaluation, 2005.
- [10] T. Funkhouser, N. Tsingos, I. Carlbom, G. Elko, M. Sondhi, J. E. West, G. Pingali, P. Min, and A. Ngan. A beam tracing method for interactive architectural acoustics. *Journal of the Acoustical Society of America*, 115(2):739–756, 2004.
- [11] M. Gander, L. Halpern, and F. Magoulès. An optimized Schwarz method with two-sided Robin transmission conditions for the Helmholtz equation. *International Journal for Numerical Methods in Fluids*, 55(2):163–175, 2007.
- [12] M. Gander, L. Halpern, F. Magoulès, and F.-X. Roux. Analysis of patch substructuring methods. *International Journal of Applied Mathematics and Computer Science*, 17(3):395–402, 2007.

- [13] M. Gander, F. Magoulès, and F. Nataf. Optimized Schwarz methods without overlap for the Helmholtz equation. *SIAM Journal on Scientific Computing*, 24(1):38–60, 2002.
- [14] S. Ghanemi. A domain decomposition method for Helmholtz scattering problems. In P. E. Bjørstad, M. Espedal, and D. Keyes, editors, *Ninth International Conference on Domain Decomposition Methods*, pages 105–112. ddm.org, 1997.
- [15] J. Günther, S. Popov, H.-P. Seidel, and P. Slusallek. Realtime ray tracing on GPU with BVH-based packet traversal. In *Proceedings of the IEEE Eurographics Symposium on Interactive Ray Tracing*, pages 113–118, Sept. 2007.
- [16] I. Harari and F. Magoulès. Numerical investigations of stabilized finite element computations for acoustics. *Wave Motion*, 39(4):339–349, 2004.
- [17] D. R. Horn, J. Sugerman, M. Houston, and P. Hanrahan. Interactive k-d tree GPU raytracing. In *Proceedings of the 2007 Symposium on Interactive 3D Graphics and Games*, pages 167–174, New York, NY, USA, 2007. ACM.
- [18] F. Ihlenburg. *Finite Element Analysis of Acoustic Scattering*. Springer, 1998.
- [19] S. Laine, S. Siltanen, T. Lokki, and L. Savioja. Accelerated beam tracing algorithm. *Applied Acoustics*, 70(1):172 – 181, 2009.
- [20] C. Lauterbach, S.-E. Yoon, M. Tang, and D. Manocha. Reducem: Interactive and memory efficient ray tracing of large models. *Computer Graphics Forum*, 27(4):1313–1321, June 2008.
- [21] Y. Maday and F. Magoulès. Non-overlapping additive Schwarz methods tuned to highly heterogeneous media. *Comptes Rendus à l’Académie des Sciences*, 341(11):701–705, 2005.
- [22] Y. Maday and F. Magoulès. Absorbing interface conditions for domain decomposition methods: a general presentation. *Computer Methods in Applied Mechanics and Engineering*, 195(29–32):3880–3900, 2006.
- [23] Y. Maday and F. Magoulès. Improved ad hoc interface conditions for Schwarz solution procedure tuned to highly heterogeneous media. *Applied Mathematical Modelling*, 30(8):731–743, 2006.
- [24] Y. Maday and F. Magoulès. Optimized Schwarz methods without overlap for highly heterogeneous media. *Computer Methods in Applied Mechanics and Engineering*, 196(8):1541–1553, 2007.
- [25] F. Magoulès. Décomposition de domaines pour le lancer de rayons. France Patent no. 1157329. 12 August 2011.
- [26] F. Magoulès, P. Iványi, and B. Topping. Convergence analysis of Schwarz methods without overlap for the Helmholtz equation. *Computers and Structures*, 82(22):1835–1847, 2004.
- [27] F. Magoulès, P. Iványi, and B. Topping. Non-overlapping Schwarz methods with optimized transmission conditions for the Helmholtz equation. *Computer Methods in Applied Mechanics and Engineering*, 193(45–47):4797–4818, 2004.
- [28] F. Magoulès, K. Meerbergen, and J.-P. Coyette. Application of a domain decomposition method with Lagrange multipliers to acoustic problems arising from the automotive industry. *Journal of Computational Acoustics*, 8(3):503–521, 2000.
- [29] F. Magoulès and R. Putanowicz. Optimal convergence of non-overlapping Schwarz methods for the Helmholtz equation. *Journal of Computational Acoustics*, 13(3):525–545, 2005.

- [30] F. Magoulès and F.-X. Roux. Lagrangian formulation of domain decomposition methods: a unified theory. *Applied Mathematical Modelling*, 30(7):593–615, 2006.
- [31] F. Magoulès, F.-X. Roux, and S. Salmon. Optimal discrete transmission conditions for a non-overlapping domain decomposition method for the Helmholtz equation. *SIAM Journal on Scientific Computing*, 25(5):1497–1515, 2004.
- [32] F. Magoulès, F.-X. Roux, and L. Series. Algebraic way to derive absorbing boundary conditions for the Helmholtz equation. *Journal of Computational Acoustics*, 13(3):433–454, 2005.
- [33] F. Magoulès, F.-X. Roux, and L. Series. Algebraic approximation of Dirichlet-to-Neumann maps for the equations of linear elasticity. *Computer Methods in Applied Mechanics and Engineering*, 195(29–32):3742–3759, 2006.
- [34] F. Magoulès, F.-X. Roux, and L. Series. Algebraic Dirichlet-to-Neumann mapping for linear elasticity problems with extreme contrasts in the coefficients. *Applied Mathematical Modelling*, 30(8):702–713, 2006.
- [35] F. Magoulès, F.-X. Roux, and L. Series. Algebraic approach to absorbing boundary conditions for the Helmholtz equation. *International Journal of Computer Mathematics*, 84(2):231–240, 2007.
- [36] R. Overbeck, R. Ramamoorthi, and W. R. Mark. A Real-time Beam Tracer with Application to Exact Soft Shadows. In *Eurographics Symposium on Rendering*, Jun 2007.
- [37] V. Pulkki, T. Lokki, and L. Savioja. Implementation and visualization of edge diffraction with image-source method. In *Proceedings of the 112th Audio Engineering Society Convention*, 2002.
- [38] F.-X. Roux, F. Magoulès, L. Series, and Y. Boubendir. Approximation of optimal interface boundary conditions for two-Lagrange multiplier FETI method. In R. K. et al, editor, *Proceedings of the 15th Int. Conf. on Domain Decomposition Methods, Berlin, Germany, Jul.21-15, 2003*, Lecture Notes in Computational Science and Engineering (LNCSE). Springer-Verlag, Heidelberg, 2005.
- [39] K. Subramanian and D. Fussel. A search structure based on k-d trees for efficient ray tracing. Technical Report Tx 78712-1188, The University of Texas at Austin, 1992.
- [40] L. L. Thompson. A review of finite-element methods for time-harmonic acoustics. *The Journal of the Acoustical Society of America*, 119(3):1315–1330, 2006.
- [41] I. Wald. *Realtime Ray Tracing and Interactive Global Illumination*. PhD thesis, Computer Graphics Group, Saarland University, 2004.
- [42] I. Wald and P. Slusallek. State of the art in interactive ray tracing. In *State of the Art Reports, EUROGRAPHICS 2001*, pages 21–42. EUROGRAPHICS, Manchester, United Kingdom, 2001.
- [43] J. Zhang, W. Zhao, and W. Zhang. Research on acoustic-structure sensitivity using FEM and BEM. *Journal of Vibration Engineering*, 18(3):366–370, 2005.