

Using volunteer computing for sound speed profile estimation in underwater acoustics*

Oleg Zaikin

Matrosov Institute for System Dynamics and Control Theory SB RAS
Irkutsk, Russia
zaikin.icc@gmail.com

Pavel Petrov

V.I. Il'ichev Pacific Oceanological Institute FEB RAS
Far Eastern Federal University
Vladivostok, Russia
petrov@poi.dvo.ru

Mikhail Posypkin

Dorodnicyn Computing Centre FRC CS RAS
Moscow, Russia
mposypkin@gmail.com

Vadim Bulavintsev

Matrosov Institute for System Dynamics and Control Theory SB RAS
Irkutsk, Russia
v.g.bulavintsev@gmail.com

Ilya Kurochkin

A.A. Kharkevich Institute for Information Transmission Problems RAS
Moscow, Russia
qurochkin@gmail.com

Abstract

In this study, we describe a volunteer computing BOINC-based project aimed at solving computationally hard inverse problems in underwater acoustics. We used this project to solve two instances of single-hydrophone dispersion-based inversion problem. This problem suits well for volunteer computing because it can be easily decomposed into independent simpler subproblems. Both instances were successfully solved, the corresponding experiments took 14 days.

*This study was partially supported by the Council for Grants of the President of the Russian Federation (grants No. MK-2262.2017.5, No. NSh-8081.2016.9), the Russian Foundation for Basic research (grants No. 16-05-01074_a, No. 16-07-00155_a, No. 15-29-07095_ofi-m, No. 17-07-00510_a and No. 16-07-00659_a), Presidium of RAS programs I.33, I.5, and the POI FEB RAS Program "Nonlinear dynamical processes in the ocean and atmosphere".

Copyright © by the paper's authors. Copying permitted for private and academic purposes.

In: E. Ivashko, A. Rumyantsev (eds.): Proceedings of the Third International Conference BOINC:FAST 2017, Petrozavodsk, Russia, August 28 - September 01, 2017, published at <http://ceur-ws.org>

1 Introduction

The notion of geoacoustic inversion refers to a collection of techniques that can be used for the reconstruction of geoacoustical waveguide structure from the sound pressure measurements [KPL12]. While normally measurements for the geoacoustic inversion are performed using expensive receiver arrays, recently it was shown that single-hydrophone recording of a broadband pulse signal can be also successfully used for estimating the acoustical parameters of sea bottom [BC11, BGNM12, BDC13]. Some promising results from a study [WDDH15] also indicate that single-hydrophone dispersion-based inversion method outlined in [BC11] can be used for the estimation of a sound-speed profile in a shallow-water waveguide. The implementation of this method in practice can be thought of as a solution of an minimization problem in a (very large) discrete search space [ZP16], and every evaluation of the cost function requires numerous solutions of an acoustic spectral problem [ZP16, Pet14]. Thus, the whole computational burden can be easily divided into a large number of relatively simple independent tasks, offering many opportunities for the application of parallel computing.

Desktop grids [CF12] (in particular, Enterprise desktop grids [IG15]) are well suited to solve hard instances of the inversion problem mentioned above. In the present paper for this purpose we use volunteer computing [Hol16] (a special kind of desktop grid computing). We launched the volunteer computing project Acoustics@home aimed at solving inverse problems in underwater acoustics. At the moment, it has the performance compared to that of small modern computing cluster. As opposed to a cluster, all project’s resources are utilized to solve inversion problems in underwater acoustics.

Let us give a brief outline of the paper. In Section 2 we overview the dispersion-based geoacoustic inversion technique. In Section 3 we describe Acoustics@home, and show the results of two computational experiments, performed in it. In Section 4 we discuss future work: new methods of black-box optimization and GPU-implementation. In the rest of the paper we discuss related work and draw conclusions.

2 Dispersion-based geoacoustic inversion

The method of the geoacoustical waveguide parameters estimation from the modal dispersion data [BC11] is rapidly gaining popularity in the underwater acoustics community [BGNM12, BDC13, WDDH15, Pet14]. By the term dispersion data we understand the set of M functions $\tau_m, m = 1, 2, \dots, M$, where $\tau_m(f)$ denotes the arrival time of m -th modal component [JKPS11, Pet14] of a pulse acoustical signal at the frequency f . The curves $t = \tau_m(f)$ in the time-frequency 2D space are called *dispersion curves* [KPL12].

The experimental dispersion curves can be obtained from a pulse signal recorded by a single receiver (hydrophone) [BGNM12, BDC13]. Normally the extraction of the dispersion curves requires some mode separation technique to be applied. A good way to separate the modal components of a pulse signal is to use the so-called warping transform [BGNM12, Boa16]. At the same time, dispersion curves can be computed theoretically provided that all the waveguide parameters are known. The mismatch between the experimental and theoretical arrival times indicates to what extent the theoretical model of a waveguide is consistent with the observation results. The set of waveguide parameters corresponding to the minimal mismatch determines the theoretical model that is the most adequate to the experimental data. We call it *media parameters estimate* based on the given experimental data. Note that an important advantage of the dispersion-based inversion schemes is their ability to provide some information on the waveguide constitution from measurements performed by a single hydrophone. By contrast, more conventional geoacoustic inversion methods [KPL12] usually require the deployment of a vertical or horizontal array of receivers (which makes the experiment much more expensive and complicated). The lack of spatial diversity of the measurements in this case is compensated for by the frequency diversity.

Thus, the geoacoustic inversion problem can be transformed into a problem of minimization of a certain mismatch function. The simplest natural choice for such function is the standard mean square fitness function measuring the average squared discrepancy between the theoretical arrivals $\tau_m^{\text{th}}(f, \mathbf{A})$ computed for the parameters vector \mathbf{A} and experimental arrivals $\tau_m^{\text{exp}}(f)$:

$$E(\mathbf{A}) = \frac{\sum_{m=1}^M \sum_{n=1}^{N_m} |\tau_m^{\text{th}}(f_{nm}, \mathbf{A}) - \tau_m^{\text{exp}}(f_{nm})|^2}{\sum_{m=1}^M N_m}. \quad (1)$$

The minimization is performed over a certain set of the admissible parameters values which is typically a cuboid in a N_p -dimensional Euclidean space, where N_p is the number of the waveguide parameters being inverted. The cuboid boundaries are determined from certain physical considerations.

In the present study we consider the geoacoustic inversion problem in the case of a homogeneous two-dimensional waveguide $\Omega = \{(x, z) | 0 \leq z \leq H\}$, where z denotes depth, and x is horizontal coordinate. The waveguide consists of the water column $0 \leq z \leq h$ and a single bottom layer $h \leq z \leq H$. The values of sound speed c_b and density ρ_b in the bottom are known constants, and we estimate the sound-speed profile (SSP) in $c = c(z)$ in the water column together with the source-receiver distance R in course of the geoacoustic inversion.

The set of the inversion parameters \mathbf{A} includes therefore R and the values of the sound speed c_1, c_2, \dots, c_{N_c} at equally spaced values of depth z_1, z_2, \dots, z_{N_c} in the water column ($z_{i+1} - z_i = \Delta z = \text{const}$). Clearly, the resolution of the SSP $c(z)$ depends on the number of nodes N_c . The more nodes we can afford, the better is the possible accuracy of the SSP estimation.

3 Computational experiments

In order to solve computationally hard problems of the kind described in Section 2, we launched the volunteer computing project Acoustics@home on 28 March 2017. The client (computing) application of this project is based on the CAMBALA MPI-program [ZP16], that we developed to solve such problems on computing clusters.

Acoustics@home is based on BOINC (Berkley Open Infrastructure for Network Computing [And04]), which is the most popular platform for volunteer computing. In Acoustics@home all daemons (that operate on the server) and computing application (it operates on volunteers' PCs) are based on CAMBALA. Work generator daemon decomposes an original problem into independent subproblems by varying several parameters of a search space. The rest parameters are varied in the computing application of the project. For each obtained set of parameters \mathbf{A} the value of the object function (1) is calculated.

We launched two experiments in Acoustics@home. The quorum [And04] of 2 was used for both of them. The first experiment was already performed on a computing cluster by our program CAMBALA in the previous work [ZP16]. We launched it in order to check the correctness of the project (in particular, the work generator and the computing applications). The corresponding input scenario no. 1 is described in Tables 1 and 2. The sound-speed profile was approximated by a piecewise-linear function with five nodes equally spaced in depth within the water column. The values of sound speed in the node points were inverted together with the source-receiver range R . The sound-speed value $c_0 = 1500$ m/s near the surface $z_0 = 0$ was assumed known (it practice it can be usually obtained from satellite data).

Table 1: Constant parameters

Parameter	value
f	50 – 300 Hz
h	50 m
H	300 m
ρ_b	1.7 g/cm ³
c_b	1700 m/s
c_0	1500 m/s

Table 2: Search space for the first scenario

Parameter	min. value	max. value	step
R	6850 m	7150 m	1 m
$c_j (j = 1, 2, 3, 4)$	1450 m/s	1500 m/s	5 m/s

According to Table 2, the corresponding search space contains 4 406 941 points (here by point we mean a set of parameters values \mathbf{A}). This search space was divided into 14 641 workunits, each of them consisted of 301 points. This experiment took 3 days, as a result we found the same global minimum (compared with that found by CAMBALA) in the corresponding search space. On average it took about 2 hours to process one such workunit on 1 CPU (Central Processing Unit) core. For the scenario no. 2 we used a piecewise-linear function with six nodes (again, equally spaced). In this case we obtained 48 476 351 points in the search space. It was divided into 161 051 workunits (301 points in each). Due to lack of resources, we could not launch this experiment on our computing cluster before. Nevertheless, Acoustics@home coped with it successfully. The corresponding experiment took 11 days. In Table 3 the obtained results for both experiments are compared with true values. In the course of the described experiments Acoustics@home had average performance of 1.5 teraflops and maximum performance of 2.5 teraflops.

Table 3: Results of geoacoustic inversion and true values.

	$E(\cdot)$	c_0 , m/s	c_1 , m/s	c_2 , m/s	c_3 , m/s	c_4 , m/s	c_5 , m/s	R , m
5 water layers								
True values	0.0221044	1500	1498	1493	1472	1462	-	7000
Scenario 1	0.0111654	1500	1465	1470	1450	1450	-	6921
6 water layers								
True values	0.0237823	1500	1499	1495	1487	1468	1462	7000
Scenario 2	0.010274	1500	1460	1475	1455	1450	1455	6922

The results of our numerical experiments are not really satisfactory from the practical point of view. Indeed, the algorithm failed to quantify the sound speed profile correctly. This can result from insufficient bandwidth of the pulse signal or from our attempt to fix the depths of the sound-speed profile nodes. In future work we will allow the latter to vary.

4 Future work

In future we plan to improve Acoustics@home in two respects. Firstly, it is intended to employ advanced black-box optimization algorithms. Secondly, a GPU-based version of the computing application will be implemented.

The problem (1) belongs to the class of black-box optimization problems which are very common in practice. In such problems derivatives are either unavailable or very hard to compute. Thus, solution methods should not rely on derivatives. In the present paper in order to minimize the objective function (1) we use uniform mesh to obtain a finite search space. This search space is processed by exhaustive search (see Section 3). In [ZP17] instead of exhaustive search we applied iterative hill climbing algorithm. In future we plan to apply the combination of global and local search techniques. The global search techniques will be used to diversify the search, by starting the local search from several initial points. The local search method will be used to minimize the value of the objective function. We plan to consider several local search techniques: Hooke-Jeeves method [HJ61], pseudo-gradient approach [ELPS16] and a variety of coordinate descent techniques.

We will also increase the performance of the computing application by utilizing GPU (Graphics Processing Unit). In comparison with CPUs, modern GPUs provide much higher computational power. Unfortunately, some algorithms could not be efficiently executed on a GPU [Bull15]. And even when there are no such obstacles, the algorithm implementation should be thoroughly planned to fit the GPU architecture. Efficient GPU implementation of an algorithm becomes possible when it demonstrates significant parallel execution opportunities, do not require random memory access, and uses no branching commands in the core computational routines.

Our approach to geoacoustic inversion problem could be represented as the following hierarchy of problems and corresponding computational algorithms (see Table 4).

Table 4: Estimated data parallelism for hierarchical levels of geoacoustic inversion algorithm.

Level	Problem	Solving method	Available parallelism
1	Inversion problem	Iterated Local Search	1-1000 (tunable)
2	Getting mismatch function for a point	Calculation of mismatch over individual frequencies	3000-10000
3	Getting modal group velocities for a frequency	Numerical differentiation	2
4	Solving Sturm-Liouville problem on a mesh	Search for eigenvalues of a tridiagonal symmetric matrix with the bisection algorithm [Dem97]	3-10

Typical implementations of the eigenvalue computation algorithms [Les07] are designed to find all eigenvalues of a matrix comprised of thousands of diagonal elements. These algorithms employ the natural "one eigenvalue per thread" strategy. However, this strategy is wasteful in our case, because we typically need only 5-10 eigenvalues distributed in a relatively narrow interval (corresponding to trapped modes). By contrast, in our implementation the parallelism is achieved by calculating the modal group velocities for thousands of frequencies simultaneously.

It should be noted that consumer-grade GPUs can't handle double precision calculations efficiently [Cor17].

Therefore, we use single precision arithmetic in GPU code. Our experiments show that usage of double precision over single precision does not accelerate convergence of ILS process and does not decrease final residue.

Table 5 shows point calculation speeds for CPU (Core i7 930, single thread) in single and double precision and for GPU (GeForce GTX 750Ti) in single precision.

Table 5: Point calculation speed for different computational platforms.

Comp. method	Performance (points/s)
CPU double precision [BB16]	0,26
CPU single precision	1,16
GPU single precision	50

5 Related work

The computing clusters are often used in the practical applications of geoacoustic inversion algorithms (see e.g. [DDH12] and references therein). While sometimes one can compromise by using heuristic optimization techniques in end-user applications, the problems of development and validation of inversion algorithms anyway set very strong demand for the high-performance computational tools.

In the previous work we have already solved the problem described in the first scenario in Section 3 on a computing cluster [ZP16]. The restrictions on available computational resources forced us to launch Acoustics@home. But we will continue to use the cluster in order to solve simple scenarios, or to test new versions of the computing application.

6 Conclusions

In the present paper we describe the volunteer computing project Acoustics@home. With the help of this project two experiments were held. The obtained results show, that this project suits well for hard instances of the considered single-hydrophone dispersion-based inversion problem. Although the present study does not offer significant contribution to the development of the geoacoustic inversion technique, it clearly illustrates the great opportunities that volunteer computing can bring into this field. In future we aim at developing a powerful computational framework that can be used on demand by any member of ocean acoustics community who needs to conduct some very intense calculations in order to solve certain direct of inverse problem. To this end, we are planning to expand our code adding some modules capable of solving propagation problems in inhomogeneous 2D and 3D waveguides.

Acknowledgments

We thank all Acoustics@home volunteers, whose computers took part in the experiment.

References

- [And04] David P. Anderson. BOINC: A system for public-resource computing and storage. In *Proceedings of the 5th IEEE/ACM International Workshop on Grid Computing, GRID'04*, pages 4–10, Washington, DC, USA, 2004. IEEE Computer Society.
- [BB16] Sergey Bochkhanov and Vladimir Bystritsky. Alglib - a cross-platform numerical analysis and data processing library. *ALGLIB Project. Novgorod, Russia*, 2016.
- [BC11] Julien Bonnel and N. Ross Chapman. Geoacoustic inversion in a dispersive waveguide using warping operators. *The Journal of the Acoustical Society of America*, 130(2):EL101–EL107, 2011.
- [BDC13] Julien Bonnel, Stan E. Dosso, and N. Ross Chapman. Bayesian geoacoustic inversion of single hydrophone light bulb data using warping dispersion analysis. *The Journal of the Acoustical Society of America*, 134(1):120–130, 2013.
- [BGNM12] J. Bonnel, C. Gervaise, B. Nicolas, and J. I. Mars. Single-receiver geoacoustic inversion using modal reversal. *The Journal of the Acoustical Society of America*, 131(1):119–128, 2012.

- [Boa16] B. Boashash. *Time-Frequency Signal Analysis and Processing (Second Edition)*. Academic Press, Oxford, 2016.
- [Bul15] Vadim Bulavintsev. An evaluation of CPU vs. GPU performance of some combinatorial algorithms for cryptoanalysis. *Vestnik Yuzhno-Ural'skogo Gosudarstvennogo Universiteta. Seriya" Vychislitel'naya Matematika i Informatika"*, 4(3):67–84, 2015.
- [CF12] Christophe Cerin and Gilles Fedak. *Desktop Grid Computing*. Chapman & Hall/CRC, 1st edition, 2012.
- [Cor17] Nvidia Corporation. Cuda c programming guide, 2017.
- [DDH12] Jan Dettmer, Stan Dosso, and Charles W. Holland. Automated geoacoustic inversion and uncertainty: Meso-scale seabed variability in shallow water environments. Report, The Office of Naval Research, 2012.
- [Dem97] James W Demmel. *Applied numerical linear algebra*. SIAM, 1997.
- [ELPS16] Yu. G. Evtushenko, S. A. Lurie, M. A. Posypkin, and Yu. O. Solyaev. Application of optimization methods for finding equilibrium states of two-dimensional crystals. *Computational Mathematics and Mathematical Physics*, 56(12):2001–2010, 2016.
- [HJ61] Robert Hooke and T. A. Jeeves. “direct search” solution of numerical and statistical problems. *J. ACM*, 8(2):212–229, April 1961.
- [Hol16] A. Holohan. *Community, Competition and Citizen Science: Voluntary Distributed Computing in a Globalized World*. Global Connections. Taylor & Francis, 2016.
- [IG15] E. Ivashko and A. Golovin. Partition algorithm for association rules mining in boinc-based enterprise desktop grid. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 9251:268–272, 2015.
- [JKPS11] F. B. Jensen, W. A. Kuperman, M. B. Porter, and H. Schmidt. *Computational ocean acoustics*. Springer, New-York et al, 2011.
- [KPL12] B. G. Katsnelson, V. G. Petnikov, and J. Lynch. *Fundamentals of Shallow Water Acoustics*. Springer US, New-York et al, 2012.
- [Les07] Christian Lessig. Eigenvalue computation with CUDA. *NVIDIA techreport*, 2007.
- [Pet14] P. S. Petrov. A method for single-hydrophone geoacoustic inversion based on the modal group velocities estimation: Application to a waveguide with inhomogeneous bottom relief. In *Proceedings of the International Conference Days on Diffraction 2014*, pages 186–191, May 2014.
- [WDDH15] Graham A. Warner, Stan E. Dosso, Jan Dettmer, and David E. Hannay. Bayesian environmental inversion of airgun modal dispersion using a single hydrophone in the chukchi sea. *The Journal of the Acoustical Society of America*, 137(6):3009–3023, 2015.
- [ZP16] Oleg Zaikin and Pavel Petrov. Algorithm of reconstruction of the sound speed profile in a shallow-water geoacoustic waveguide from modal dispersion data. *Optoelectronics, Instrumentation and Data Processing*, 52(3):259–265, 2016.
- [ZP17] Oleg Zaikin and Pavel Petrov. Application of iterative hill climbing to the sound speed profile inversion in underwater acoustics. In *Proc. of The Fifth International Workshop on Mathematical Models and their Application*, volume 173 of *IOP Conf. Series: Materials Science and Engineering*, pages 1–8, 2017.