Acoustic Propagation Modelling on the Cloud

R. Marques Rojo, J.D. Gonzalez, S. Blanc

Argentinian Navy Research Office & UNIDEF (CONICET)

INTRODUCTION AND MAIN RESULTS

The **computation of 3D acoustic models for long-range propagation in realistic oceanic environments** poses significant computational challenges. Here, **we present an extension of our GPU-accelerated hydroacoustic transmission loss solver in Julia.** This work is built upon a model based on the parabolic wave equation and the **3D Split Step Fourier method**, we study its capability **to run in cloud environments with large-memory GPU instances**. The extended solver can resume long computations that were interrupted, enabling it to leverage interruptible (spot) nodes thus reducing computational costs. Preliminary tests confirm that this approach indeed enables resolutions, swath widths, frequencies, and propagation distances previously unattainable on our standard desktop systems, **significantly reducing the need for substantial capital investments in dedicated HPC resources**. We illustrate this with a case study involving a propagation distance of 10,000 km. This work is intended to lower barriers to large-scale hydroacoustic research by enabling analyses of more detailed scenarios and providing a flexible and cost-effective solution for diverse applications in underwater acoustics research, potentially aiding in aspects such as the detection and analysis of underwater acoustic events like those **monitored by the IMS**.

SnT 2025
8 SEPTEMBER ONLINE DAY
9 TO 12 SEPTEMBER AT HOFBURG PALACE, VIENNA & ONLINE
CTBT: SCIENCE AND TECHNOLOGY CONFERENCE

Acoustic Propagation Modelling on the Cloud

R. Marques Rojo, J.D. Gonzalez, S. Blanc

P1.3-766

Introduction

- **Problem**: Realistic long-range (10^3–10^4 km) hydroacoustic TL needs 3D effects (island shielding/diffraction, bathymetry-coupled refraction) and is bottlenecked by GPU VRAM and wall-time on desktops.
- **Code**: Our team has been implementing the SplitStepFourier.jl *Julia* package that efficiently computes acoustic transmission losses in the water and sediment column on consumer grade CPU or GPU hardware.
- **Model:** 3D Parabolic Equation with Split-Step Fourier, marching in range.
- **Data pipeline:** automatic fetch as needed:
  - Bathymetry from GEBCO
  - Temperature/salinity profiles from Copernicus reanalysis db.

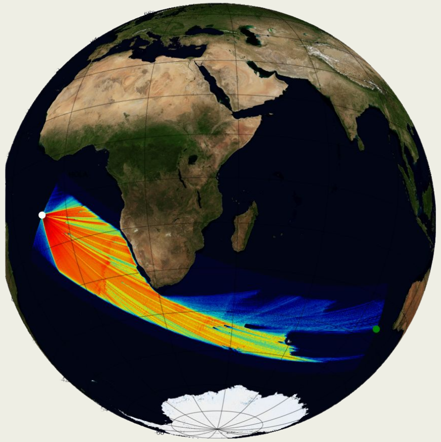

Key inputs

- Geographic positions
- Date
- Frequency
- Source depth
- Swath width

Implementation

- Julia + CUDA (optional).
  - Array-generic; CUDA via dispatch.
- Recent advances:
  - **Checkpoint/resume** (hash-based run ID + regular state saved to disk) to immunize from interruption risks in more economical Cloud spot instances.
  - **Fused FFT**s; asynchronous post-processing; fewer kernel launches; environment precompute reused during frequency sweeps; reduced allocations; among others.
- 2-5x speedup from previous implementation.
- Observed 10-100x speedup vs multi-threaded run on Intel Xeon(R) Gold 6240R@2.40GHz with 48 physical cores and 256 GB RAM. Exact speedup depends on GPU and simulation specific settings.

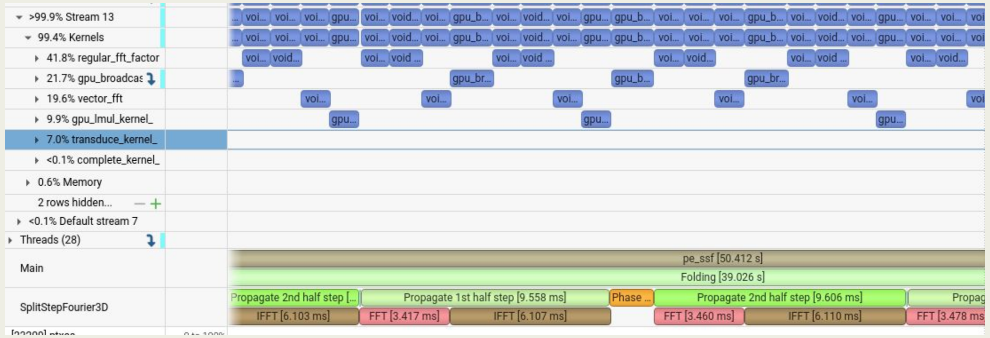

Capture of Nsight Systems CUDA profiler during the marching steps

Illustrative scenario (South Atlantic crossing from H10N)

3D run in Cloud based GPU (top) vs Nx2D RAM local run (bottom) @ 20 Hz

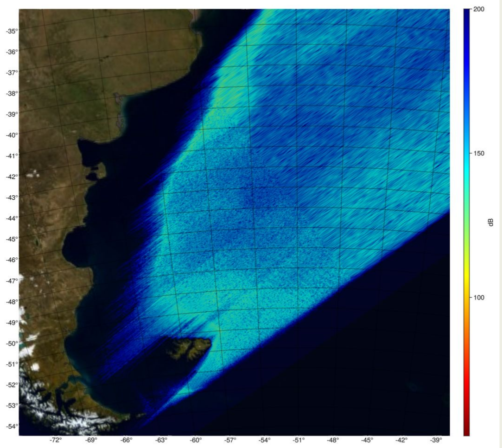

3D

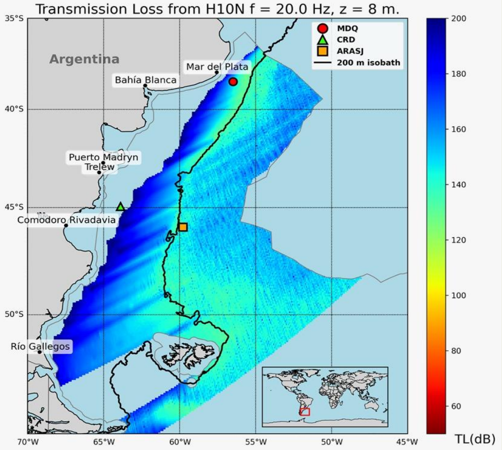

Nx2D

CTBTO PREPARATORY COMMISSION

SnT 2025
CTBT: SCIENCE AND TECHNOLOGY CONFERENCE
8 SEPTEMBER ONLINE DAY
9 TO 12 SEPTEMBER AT HOFBURG PALACE, VIENNA & ONLINE

Acoustic Propagation Modelling on the Cloud

R. Marques Rojo, J.D. Gonzalez, S. Blanc

Results I

- Simulation run at f = 2 kHz, with a propagation range of 7,420 km, and Δz fixed at 8 m.
- Plot shows TL (dB) at z = 48 m.
- Diffraction effects around the Malvinas Islands are observed, which cannot be reproduced by Nx2D models.

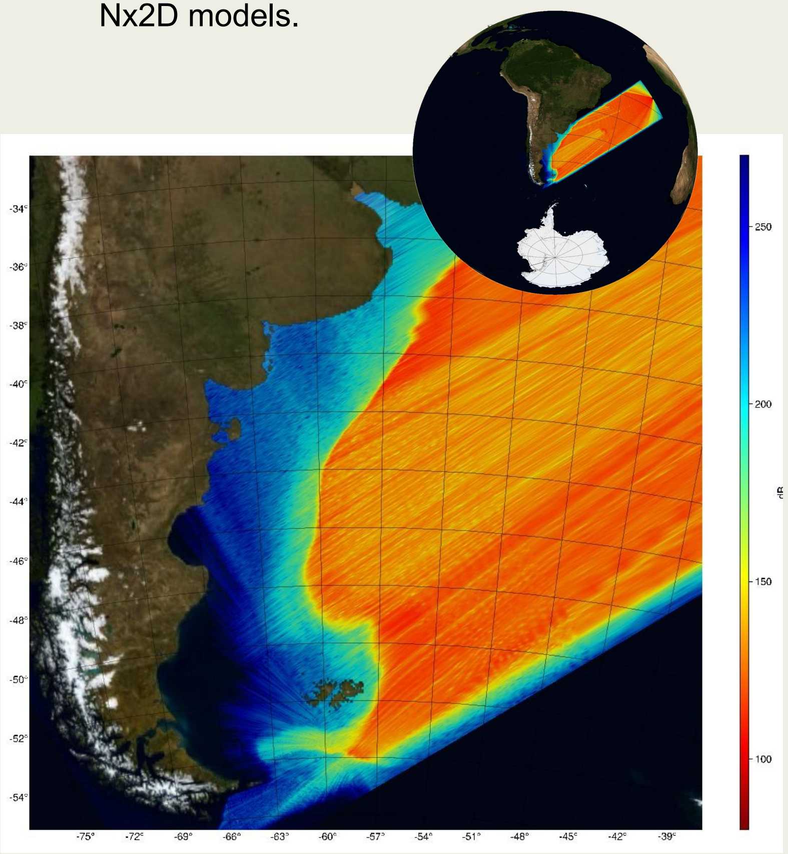

Results II

The previous scenario was evaluated at multiple frequencies, with grid steps defined as $\Delta x = \lambda/4$, $\Delta y = \lambda/2$, and $\Delta z = \lambda/10$;
Swath width ≈ 2000 km, z maximum ≈ 11 km
Three execution times obtained on a cloud-based GPU are reported in the following Table.

Frequency (Hz)	Screen nodes (n_z x n_y)	Runtime
5	10.6 M	5 min 15 s
10	40.3 M	30 min 54 s
16	107.5 M	2h 01 min 7 s

- These timings correspond to runs on Google Cloud A3 High (*a3-highgpu-1g*) instances with 1x NVIDIA H100 SXM GPU.
- Spot instance in *us-east4* region at **$2.04/hr** (regular on-demand price is $11.07/hr) - **82% discount** (*compute cost only*).
- A separate run at 28 Hz, with **229.5M screen nodes**, ran **under 11 GiB VRAM**.
  - We **could not have run that locally** with our available GPU at the time.
  - There is still much more room to grow the problem size with cloud resources, even before requiring multi-node multi-GPU setups.

Conclusion

- **3D models** better capture **diffraction** effects, offering more realistic insights into underwater acoustic propagation.
- Even **array-generic code**, with no special CUDA tuning, **can deliver huge speedups**.
- **Cloud computing** provides access to high-performance resources at relatively **low cost**, making this approach particularly valuable for small institutions and research groups in low-income countries.
- It lets you **try before you buy**.
- **Spot pricing** cut cost a lot; **checkpointing** is key for long runs (in practice we saw only one interruption over many hours of computation).
- **Do short tests before long marches** to reduce the risk of incurring in unexpected costs.

CTBTO
PREPARATORY COMMISSION