

NETSIM: A Realtime Virtual Ocean Hardware-in-the-loop Acoustic Modem Network Simulator

Toby Schneider
GobySoft, LLC
Woods Hole, MA, USA
toby@gobysoft.org

Henrik Schmidt
Department of Mechanical Engineering
Massachusetts Institute of Technology
Cambridge, MA, USA
henrik@mit.edu

Abstract—This paper presents *netsim*, a combined software/hardware system for performing realtime realistic operation of autonomous underwater vehicles (AUVs) with acoustic modem telemetry in a virtual ocean environment. The design of the system is flexible to the choice of physical link hardware, allowing for the system to be tested against existing and new modems. Additionally, the virtual ocean channel simulator is designed to perform in real time by coupling less frequent asynchronous queries to high-fidelity models of the ocean environment and acoustic propagation with frequent perturbation-based updates for the exact position of the simulated AUVs. The results demonstrate the performance of this system using the WHOI Micro-Modem 2 hardware in the virtual ocean environment of the Arctic Beaufort Sea around 73 degrees latitude. The acoustic environment in this area has changed dramatically in recent years due to the changing climate.

Index Terms—simulation, virtual ocean, hardware-in-the-loop, acoustic communication, undersea networking

I. INTRODUCTION

The development of successful autonomous underwater vehicle (AUV) systems requires substantial simulation capability due to the high cost and low availability of at-sea testing time. As the components to be simulated and the physics of the environment with which the components interact increases in complexity, it becomes more difficult to develop simulators with a reasonable degree of accuracy. Acoustic communication subsystems are such a component: modeling of acoustic propagation for acoustic data carriers is a rich area of current research as is the related discipline of the design of acoustic modem systems and signal processing. To further complicate the problem for marine roboticists, many acoustic modems that are mature enough to be fielded are at least partially proprietary in design.

Thus, the paper presents a software/hardware hybrid system (called *netsim*) for realtime simulation of acoustic modem signals in a virtual ocean. The virtual ocean model it presents is based around asynchronous calls to a 4-dimensional ocean model (MSEAS) and to a raytracing model (BELLHOP). These high-fidelity, but relatively computationally intensive

results are updated synchronously with the precise simulation positions of the AUVs using perturbations around the asynchronous results. This approach provides realtime performance (at ranges above a few hundred meters).

The results section of the paper studies a case that is of immediate concern to AUV roboticists: the decreased performance of sonar and modem systems in the Beaufort Sea due to climate change effects. The discovery and background of this problem is presented in [1], [2].

A. Prior work

Numerous channel simulators have been developed for underwater acoustics (e.g. [3], [4]), many of which have been presented at previous UComms workshops [5]. This work is not intended to supplant these, but rather provide a motivation for creating channel simulators that can be directly applied to robotic systems.

Similarly, numerous acoustic modems exist for both research and fielded applications, such as the WHOI Micro-Modem, the Teledyne Benthos modem, the EvoLogics modems, and others.

Peterson and Porter [6] presented a partial hardware-in-the-loop study. They used playback data captured from the acoustic modem to characterize the performance of their channel simulator. This work provides simulation in real time so that the acoustic modem performance can be simulated simultaneously with the rest of a realtime AUV system.

In addition, several *in-situ* hardware testbeds have been developed, such as those in [7] and [8]. These in-water testbeds are a highly valuable resource, but are limited to the physical layout of the particular testbed deployment, and often have limited support for remote testing of mobile nodes, due to the practical concerns involved in deploying and recovering AUVs. Also, the real ocean provides very limited reproducibility; therefore, in-situ testing at a given location does not allow for a comprehensive comparison of different modem technologies and signal processing approaches.

II. DESIGN

One side of the *netsim* system handles analog input and output (I/O) to acoustic modem hardware, whereas the other

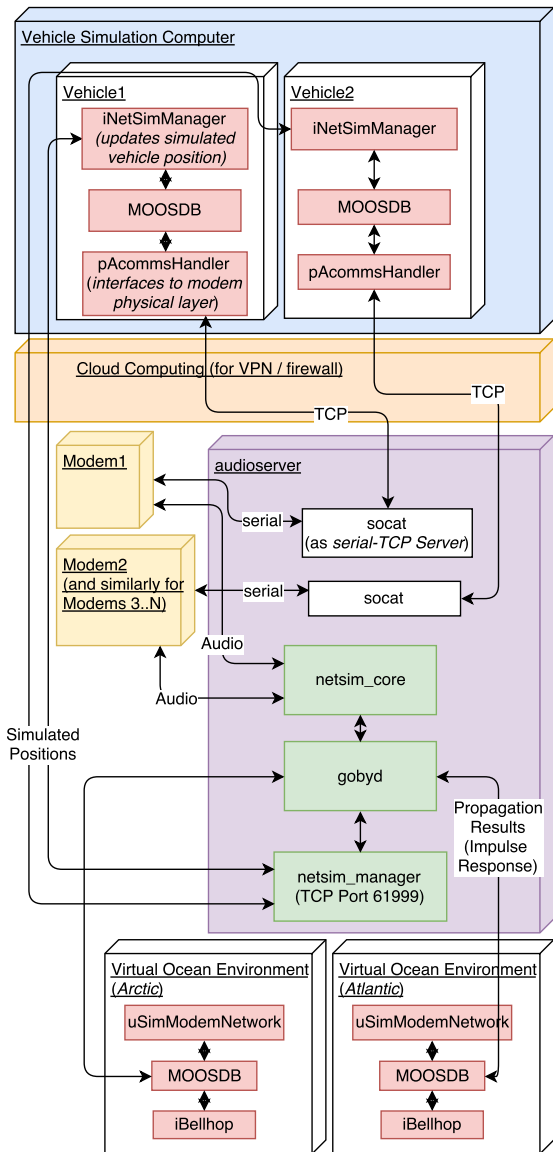


Fig. 1. The *netsim* system showing both hardware and software components. Red components are written in MOOS [9], green are in Goby3 [10]. The vehicle simulation computer can be located anywhere with an Internet connection; the modems are co-located with the *audioserver* due to the high-bandwidth/low-latency requirements of the audio.

side handles interfacing to a oceanographic model combined with a channel simulator (the “Virtual Ocean”). These two ends are flexible: any modem sampled at 96kHz or below can be used to feed the modem I/O, and any channel simulator capable of performing within the realtime requirements can be coupled to the other side.

Fig. 1 diagrams the entire *netsim* system. For the flexibility of the users, the “Vehicle Simulation Computer” can be located anywhere in the world, with an Internet connection capable of a few kB/s throughput or better, since all that is transferred over this link is acoustic modem traffic and regular (e.g. 1 Hz) simulated vehicle position updates.

The *audioserver* comprises the audio processing compo-

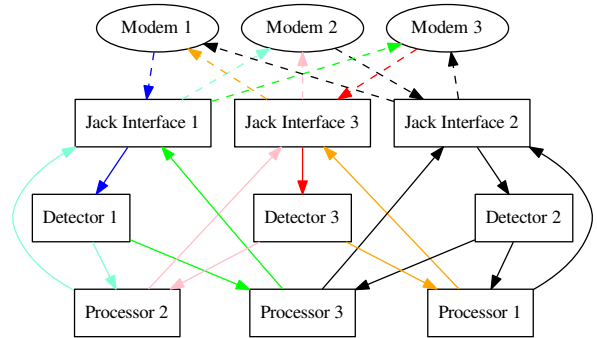


Fig. 2. Data flow for *netsim-core* showing three modems. Dashed lines represent analog audio; solid lines represent Goby3 digital audio frames. The colors represent two possible transmission paths: 1) blue is a transmission from modem 1 which becomes teal and green receptions for modem 2 and 3, respectively. 2) red is a transmission from modem 3, which becomes orange and pink receptions for modem 1 and 2. Depending on temporal/spatial separation of the simulated modem locations and transmission times, collisions may occur and are modeled in *netsim* by audio mixing at modem 2.

nents and physical (analog audio) interface to the acoustic modems. The “Virtual Ocean Environments” can be co-located with the *audioserver* or over a local TCP connection. The modular design of the system intends to place minimal restrictions on the modem hardware used or on the Virtual Ocean specifics. This will allow testing of numerous modem designs, along with the continued improved fidelity of the Virtual Ocean.

A. Realtime audio processing

A key part of successfully performing realtime audio simulation is providing low-latency processing and manipulation of the audio frames (where a frame is defined as a set of N samples, where $N=4096$ for the results presented in section IV).

The realtime audio processing was performed by combining the Jack audio toolkit on Linux [11] with the interthread layer of the C++ Goby3 publish/subscribe middleware [10] into an application called *netsim-core*. Each modem is assigned an input and output channel and mixing at the receiver is done by Jack. All the remaining audio frame handling is done by three groups of threads within Goby3, which are diagrammed in Fig. 2:

- 1) Jack Interface: Interface threads to the Jack toolkit which interfaces with the analog I/O via the Linux audio system.
- 2) Detector threads: Detects start and end of a packet for logging purposes.
- 3) Processor threads: Runs the channel simulator for each packet (on a frame basis).

III. THE VIRTUAL OCEAN AUTONOMY TESTBED

The Virtual Ocean MOOS-IvP Autonomy Testbed illustrated in Fig. 3 replicates the physical ocean environment

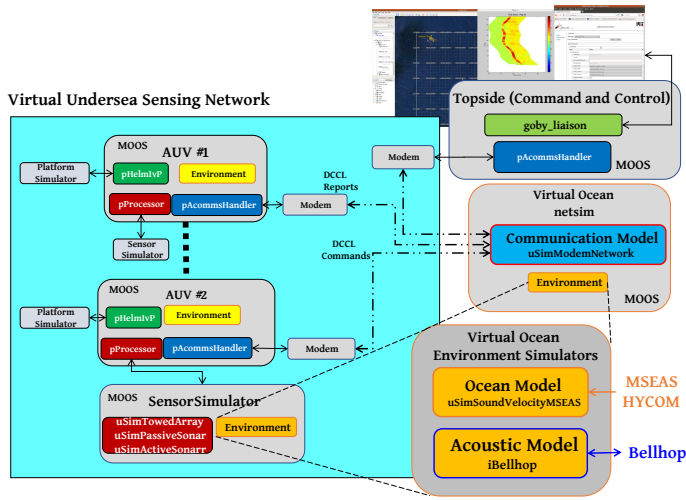


Fig. 3. **Virtual Ocean Autonomy Testbed:** The unaltered Platform Autonomy System is running missions with the three connections to the real ocean, the sensor, the modem and the platform, replaced by simulated versions connected to a high-fidelity, physics-based environmental ocean and acoustic modeling and forecasting infrastructure.

surrounding an autonomous underwater vehicle or mooring using a high-fidelity, physics-based environmental modeling infrastructure, the *Virtual Ocean*. The architecture is designed to make the MOOS-IvP Platform Autonomy system, including the decision-making Helm *pHelmIvP*, agnostic to whether it is deployed in a real or virtual ocean. This is accomplished by replacing each of the three physical connections of each payload autonomy system, the low-level platform control, the sensors, and the modem with high-fidelity, physics-based simulators. By building the simulators as separate MOOS communities, the modularity allows for the simulators to be incrementally developed and implemented as the missions demand. At the center of the *Virtual Ocean* is the representation of the physical properties of the ocean in the operational area by a 4-D (time, latitude, longitude and depth) database of current, temperature, salinity, and bathymetry, typically produced by a modeling and assimilation framework such as MSEAS [12] or HYCOM [13] or via historically sampled data. This common environmental picture supports the various physical models of the “outside world” for the autonomy system. Thus, the model of the platform dynamics uses the current field, and the simulators for oceanographic properties sample this data matrix in space and time, using the same interface to the autonomy system as the equivalent field sensors. Similarly, all acoustic sensor, communication and navigation systems extract environmental sound speed, bathymetry etc. from the database for driving the computational acoustic modules. Also here, the interface to the onboard autonomy system is fully replicated by the environmental acoustic simulators producing element-level timeseries of the signal and noise, directly stimulating the full “runtime” or fielded version of the acquisition and processing chains.

A snapshot of the *Virtual Ocean* is embedded on the

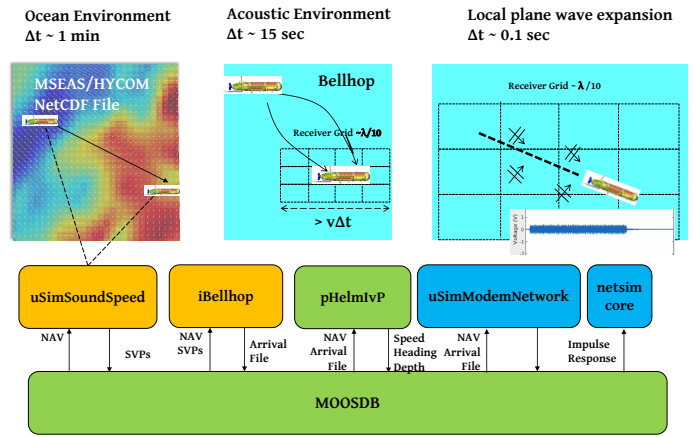


Fig. 4. Nested modeling and simulation infrastructure for real-time element-level acoustic stimulation of sensors and modems. In the embedded version, the Helm and other processes request transfer function or TL/SNR estimates for performance optimization, while the timeseries stimulators like *netsim* requests the impulse response and synthesizes element-level timeseries fed to the sensor acquisition or the incoming modem.

network nodes, where it supports the onboard data processing, mission planning and decision making. To accurately replicate real ocean operations, the platform version must be made ignorant of the dynamics of the ground-truth version used by the simulators. This is accomplished by having messages to update the simulators that are out-of-band from the mission operational data.

A. Timescale driven realtime network simulation

In the context of *netsim*, the relevant *Virtual Ocean* component is the real-time acoustic array simulator. Providing high-fidelity simulation in real time of element-level timeseries of incoming modem or acoustic signals between dynamic platforms, the key enabler is the nested environmental acoustic modeling and sensor simulation framework illustrated in Fig. 4. The nesting is achieved by heavy use of pre-computing (asynchronously with vehicle position updates) of environmental features consistent with their timescales of change. Linked to legacy ocean modeling and assimilation frameworks such as MSEAS and HYCOM, the simulation framework supports realistic dynamic changes in the environment. This allows for efficient performance assessment of environmental adaptation autonomy. However, the underlying ocean representation needs to be updated only at intervals consistent with the slowly varying ocean. Similarly, the use of legacy acoustic models with decomposable signal components such as rays (BELLHOP) or modes (Kraken) limits the computational updates to be required on timescales consistent with the macroscopic changes in source/receiver configuration. Thus, a process requiring an acoustic impulse response between two platforms, such as the impulse-response engine *uSimModemNetwork*, will request the sound speed distribution between the two nodes every few minutes and request the impulse response from the acoustic model every 15-30 seconds for a grid covering the receiver dynamics over the same period.




Fig. 5. For a transmitter at 33 m (a) and 100 m (b) depth in the simulated Arctic environment centered at 73.128° N, 149.400° W, the mean squared error (MSE) of Micro-Modem [15] equalizer for transmissions to a grid of simulated receiver locations through *netsim*. For each of these grid points, a 72-byte Micro-Modem Rate 1 “Flexible Data Packet” was transmitted with carrier frequency of 3500 Hz and a bandwidth of 1250 Hz. These grid points are overlaid on an incoherent transmission loss rendering using BELLHOP [16].

Incorporating a multi-path Doppler framework [14], a simple plane wave interpolation within this receiver grid then provides a high fidelity element-level impulse response on an arbitrary volumetric array on the moving receiving platform, capturing all the spatial doppler effects. The received timeseries are then produced in short time, typically 0.1 second bursts, by convolving each individual multipath impulse with the associated Doppler-shifted source replica. Short timescale environmental variability is incorporated using stochastic models for the signal phase, including legacy surface scattering models and stochastic source models.

Supporting simulation of realistic ambient noise on the receiving arrays, a similar nested modeling infrastructure uses legacy models for surface generated noise covariance to provide the slowly varying vertical noise directionality, which is decomposed into a spherical harmonic representation. This representation in combination with the current spectral noise level is efficiently translated into a noise covariance for an arbitrary dynamic, volumetric array configuration. Choleski factorization of the resulting covariance matrix is then used for generating the array element timeseries consistent with the dynamic noise correlation properties [14].

IV. RESULTS

We applied *netsim* to several environments, but will present the Arctic environment in the Beaufort Sea as it is a compelling case for vehicle adaptivity and the need for accurate acoustic propagation modeling applied to robotics systems. For this test, *netsim* was capable of providing realtime performance for eight modems on a Intel Xeon X5450 (8-core 3.0 GHz). This includes both the *audioserver* and Virtual Ocean components. In the future we will likely move the *audioserver* (which has a hard realtime constraint) to a dedicated lower-

power CPU to separate the realtime constraints into separate processors. This should allow simulation of more modems without missing any realtime deadlines on the *audioserver*.

A. Historical transducer placement: 33 meters deep

Fig. 5a models an transmitter (with 120 degree beamwidth) located at 33 meters depth, which was used in the ICEX16 sea trial. This depth historically was sufficient for AUV operations around 100m depth, but the new Beaufort lens effect [2] creates a shadow zone below 50 meters starting at about 2 km. To demonstrate this effect on modem performance, we ran 1500 transmissions through the Micro-Modem 2, using the exact settings used in ICEX16, varying the receiver location over a fixed grid. These results were color-coded based on whether the packet was received (no reception: red) and if so, the error measured by the Micro-Modem equalizer (high error: blue, low error: green). For reference, we overlaid the results of these transmissions on a transmission loss plot using BELLHOP’s finite element ray tracing technique.

Not surprisingly, the shadow zone regions showed no packets or high equalizer error, and the regions with lower transmission loss tend to show valid packet reception. However, the transmission loss does not show the whole picture since acoustic modem performance is coherent. Regions with the best packet reception tend to be a low number of rays with a clear direct path, such as the caustic that runs along the edge of the shadow zone from around $(r, z) = (1000, -50)$ to $(3500, -250)$ meters.

B. Potential new transducer placement: 100 meters deep

The modeling, backed by the results of using the full hardware-in-the-loop *netsim*, suggests that a transducer placed at 100 meters will provide more uniform communications performance in the new Arctic environment. However, this may not always be operationally feasible, which is the motivation behind adaptive behaviors for improving communications performance, such as those presented in [17]. The *netsim* system allows us to robustly test such behaviors in a high-fidelity simulated environment.

ACKNOWLEDGMENT

We thank the developers of the open source software components used in this work: the Jack Audio team, Paul Newman (MOOS), Mike Benjamin (MOOS-IvP), and Mike Porter (BELLHOP). In addition, we thank all those involved in making the ICEX16 experiment possible, from which we were inspired to find new solutions. The authors appreciate the effort of LT Scott Carper, Thom Howe, and Andrew Poulsen, without which the ICEX16 AUV deployment, which inspired and motivated this work, would never have been successfully completed. We also appreciate the tireless effort by Dr. Bob Headrick of ONR Code 322OA on securing funding for the MIT ICEX16 effort from the ONR Experimentation program lead by Dr. Tom Killian. The support from USN DEVRON-12, Art Baggeroer MIT, and ADM Mike Connor in making the ICEX16 deployment happen in the first place is highly appreciated.

REFERENCES

- [1] L. Freitag, K. Ball, J. Partan, P. Koski, and S. Singh, "Long range acoustic communications and navigation in the arctic," in *OCEANS'15 MTS/IEEE Washington*. IEEE, 2015, pp. 1–5.
- [2] H. Schmidt and T. Schneider, "Acoustic communication and navigation in the new arctic model case for environmental adaptation," in *Underwater Communications and Networking Conference (UComms), 2016 IEEE Third*. IEEE, 2016, pp. 1–4.
- [3] P. Qarabaqi and M. Stojanovic, "Statistical characterization and computationally efficient modeling of a class of underwater acoustic communication channels," *IEEE Journal of Oceanic Engineering*, vol. 38, no. 4, pp. 701–717, 2013.
- [4] M. S. Caley and A. J. Duncan, "Wide-band shallow acoustic channel simulation with realistic Doppler and delay spreading for 3D evolving rough surfaces," in *2016 IEEE Third Underwater Communications and Networking Conference (UComms)*, Aug 2016, pp. 1–5.
- [5] J. R. Potter, M. B. Porter, and J. C. Preisig, "UComms: A conference and workshop on underwater communications, channel modeling, and validation," *IEEE Journal of Oceanic Engineering*, vol. 38, no. 4, pp. 603–613, Oct 2013.
- [6] J. C. Peterson and M. B. Porter, "Ray/beam tracing for modeling the effects of ocean and platform dynamics," *IEEE Journal of Oceanic Engineering*, vol. 38, no. 4, pp. 655–665, 2013.
- [7] R. Martins, J. B. de Sousa, R. Caldas, C. Petrioli, and J. Potter, "SUNRISE project: Porto university testbed," in *2014 Underwater Communications and Networking (UComms)*, Sept 2014, pp. 1–5.
- [8] J. Alves, J. Potter, P. Guerrini, G. Zappa, and K. LePage, "The loon in 2014: Test bed description," in *2014 Underwater Communications and Networking (UComms)*, Sept 2014, pp. 1–4.
- [9] M. R. Benjamin, J. J. Leonard, H. Schmidt, and P. M. Newman, "An overview of moos-ivp and a brief users guide to the ivp helm autonomy software," 2009.
- [10] T. Schneider, "Goby3: A new open-source middleware for nested communication on autonomous marine vehicles," in *Autonomous Underwater Vehicles (AUV), 2016 IEEE/OES*. IEEE, 2016, pp. 236–240.
- [11] S. Letz, D. Fober, and Y. Orlarey, "Jack audio server for multi-processor machines," in *ICMC*. Citeseer, 2005.
- [12] P. Lermusiaux. Multidisciplinary simulation, estimation, and assimilation systems. [Online]. Available: <http://mseas.mit.edu>
- [13] Hybrid coordinate ocean model. [Online]. Available: <https://hycom.org/>
- [14] F. Jensen, W. Kuperman, M. Porter, and S. H., *Computational Ocean Acoustics*. New York, NY: Springer, 2011.
- [15] E. Gallimore, J. Partan, I. Vaughn, S. Singh, J. Shusta, and L. Freitag, "The WHOI Micromodem-2: A scalable system for acoustic communications and networking," in *OCEANS 2010 MTS/IEEE SEATTLE*, Sept 2010, pp. 1–7.
- [16] M. B. Porter, "The BELLHOP manual and users guide: Preliminary draft," *Heat, Light, and Sound Research, Inc., La Jolla, CA, USA, Tech. Rep*, 2011.
- [17] T. Schneider and H. Schmidt, "Model-based adaptive behavior framework for optimal acoustic communication and sensing by marine robots," *IEEE Journal of Oceanic Engineering*, vol. 38, no. 3, pp. 522–533, 2013.