

An architecture for passive joint localization and structure learning in reverberant environments

Toros Arikan,^{1,a)} Amir Weiss,¹ Hari Vishnu,² Grant B. Deane,³ Andrew C. Singer,⁴ and Gregory W. Wornell¹

¹Electrical Engineering and Computer Science Department, Massachusetts Institute of Technology, Cambridge, Massachusetts 02142, USA

²Acoustic Research Laboratory, Tropical Marine Science Institute, National University of Singapore, Singapore 119222

³Scripps Institution of Oceanography, University of California San Diego, La Jolla, California 92037, USA

⁴University of Illinois at Urbana Champaign, Urbana, Illinois 61820, USA

ABSTRACT:

Passive localization and tracking of a mobile emitter, and joint learning of its reverberant three-dimensional (3D) acoustic environment, where critical structural features are unknown, is a key open problem. Unaccounted-for occluders are potentially present, so that the emitter can lose line-of-sight to the receivers, and can only be observed through its reflected raypaths. The locations of reflective boundaries must therefore be jointly estimated with the emitter's position. A multistage global optimization and tracking architecture is developed to solve this problem with a relatively unconstrained model. Each stage of this architecture establishes domain knowledge such as synchronization and initial environment estimation, which are inputs for the following stages of more refined algorithms. This approach is generalizable to different physical scales and modalities and improves on methods that do not exploit the motion of the emitter. In one stage of this architecture, particle swarm optimization is used to simultaneously estimate the environment and the emitter location. In another stage, a Hough transform-inspired boundary localization algorithm is extended to 3D settings, to establish an initial estimate of the environment. The performance of this holistic approach is analyzed and its reliability is demonstrated in a reverberant watertank testbed, which models the shallow-water underwater acoustic setting. © 2023 Acoustical Society of America.

<https://doi.org/10.1121/10.0016999>

(Received 11 August 2022; revised 15 December 2022; accepted 6 January 2023; published online 27 January 2023)

[Editor: Stan E. Dosso]

Pages: 665–677

I. INTRODUCTION

Passive underwater acoustic localization in shallow-water settings is one of the most important and challenging localization problems (Fischell *et al.*, 2019; Waterston *et al.*, 2019). The emitter to be localized could be a surface vessel or underwater vehicle, and the receivers could be surreptitiously placed for passively monitoring a body of water. Receivers are typically omnidirectional tethered hydrophones, but could also be placed on a set of mobile underwater vehicles, as illustrated in Fig. 1. When the environment is accurately known, state-of-the-art matched-field processing (MFP) methods can be used for localization (Baggeroer *et al.*, 1993). The standard MFP framework does not account for the presence of man-made features such as pier pilings or anchored vessels, or unsurveyed reefs (Chitre and Pelekanakis, 2014). Depending on their positions, these large and potentially unknown objects can act as additional reflective boundaries, or as occluders that block the line-of-sight (LOS) arrivals to receivers. MFP is sensitive to model mismatch (Le Gall *et al.*, 2016; Xu *et al.*, 2006) and its performance could therefore deteriorate due to such objects,

motivating the formulation of a localization framework that incorporates their presence. A representative scenario is presented in Fig. 2, where the goal is to estimate the unknown source position at each transmission point along a trajectory (red circles), in the presence of a large occluder, in a scaled model of the shallow-water setting. This is an instance of an open problem that also arises for indoor acoustic or radio frequency localization (Dokmanic *et al.*, 2015), where the non-line-of-sight (NLOS) arrivals due to multipath reflections must be judiciously exploited rather than mitigated, with a robust passive localization method.

Without restrictions on the localization setting, such problems are ill-posed and admit infinitely many solutions. The static nature of the environment, and the spatial continuity in the trajectory of a moving emitter, are key to unambiguous localization. Some of the important aspects of the setting are modeled as in Table I; the limited information here leaves some key parameters undetermined so that it is natural to adopt a multistage approach to first estimate them. Accordingly, we deploy a series of localization algorithms sequentially, as depicted in Fig. 3, using an architecture that we call passive end-to-end localization (PEEL). Each stage of PEEL produces approximate solutions to different aspects

^{a)}Electronic mail: atoros@mit.edu

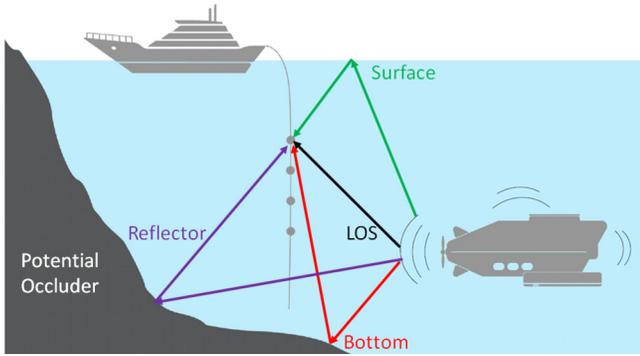


FIG. 1. (Color online) Typical shallow-water underwater acoustic localization problem setting [image is modified from the original in Kim *et al.* (2019); licensed under a Creative Commons Attribution 4.0 License].

of the localization problem, which then become inputs for the following stages. We call this an end-to-end method because it starts from received waveforms that do not have a common clock with the emitter, and delivers a joint estimate of the full emitter trajectory and the environment. We ultimately demonstrate that PEEL achieves strong performance, even when there is modeling mismatch that violates the assumptions in Table I.

While there is currently no end-to-end passive localization method such as the one we propose, similar frameworks have been designed for different settings. Joint acoustic tracking and environment learning problems can be solved using simultaneous localization and mapping (SLAM) methods (Li and Krolik, 2013; Ribeiro *et al.*, 2010) in active settings, or when the emitter and receivers form a network (Naseri and Koivunen, 2016). Similarly, while there are time-difference-of-arrival (TDOA) and time-of-arrival (TOA) methods for localization (Scheuing and Yang, 2008; Venkateswaran and Madhow, 2012; Alameda-Pineda and Horaud, 2014), they do not account for unknown occluders as PEEL does.

Our main contribution is an architecture for joint localization and environment learning. Among other tools, PEEL makes use of:

TABLE I. Model of the localization setting.

Problem feature	Modeling assumptions
Speed of sound, v_s	Known, and constant within the environment
Environment	Static; only the emitter position is changing
Reflectors, η_i	Planar, known number, unknown position and orientation
Transmissions	Periodic transmissions of deterministic pulses, with known period
Receivers, $\mathbf{p}_{r,j}$	At least 4 receivers deployed, with known positions

- A global optimization method for joint boundary estimation and localization, which does *not* need to be tuned to a specific problem.
- The mapping of unknown occluders after tracking, followed by iterations with a refined optimization metric, which has not been implemented for passive localization.
- Closed-form derivations for some key steps of a boundary localization technique (Naseri and Koivunen, 2016), that improve computational efficiency and accuracy.

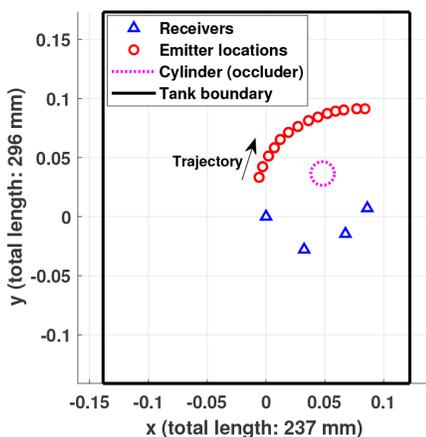
The paper is structured as follows. In Sec. II, we present the signal model and geometric considerations used for localization. We present the PEEL method in Sec. III. The results that are achieved in the scale watertank environment of Fig. 2 are given in Sec. IV. In Sec. V, we conclude with potential avenues for future research.

In our mathematical notation throughout the paper, lowercase bold variables such as \mathbf{p} are vectors, and uppercase calligraphic characters such as \mathcal{P} denote sets.

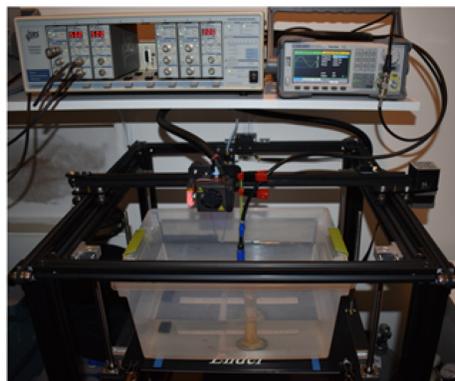
II. LOCALIZATION PRINCIPLE: EXPLOITING VIRTUAL RECEIVERS

In this section, we introduce the notation and signal model, and present key geometrical considerations that are necessary to frame the localization problem. We specifically focus on how multipath arrivals can be modeled and used to improve localization performance.

Triangulation is a standard localization principle that models the emitter as the point of intersection of circles of equal time-of-flight (TOF) from the emitter to each receiver.



(a)



(b)

FIG. 2. (Color online) Top-down view of the test environment with emitter's trajectory (a), and the actual experimental setup (b). This is a 3D localization and tracking problem with an occluding cylinder, and is more challenging than settings where the receivers surround the emitter (Setlur *et al.*, 2012).

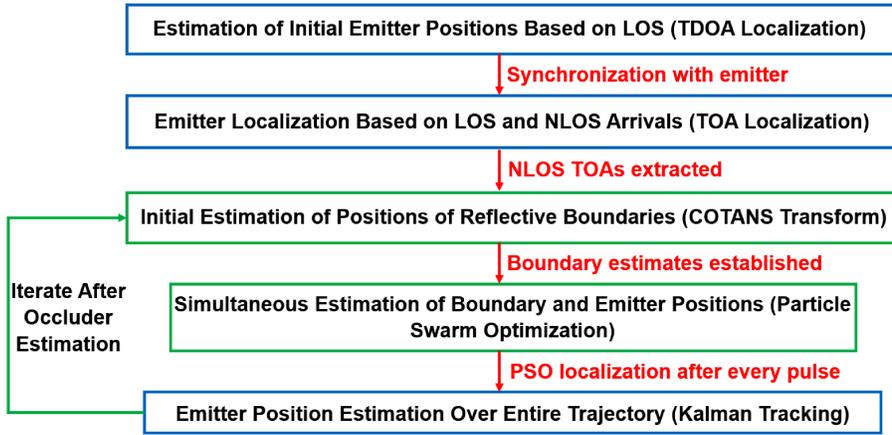


FIG. 3. (Color online) The sequence of algorithms in the PEEL method. The sections with green outlines incorporate novel developments or applications of those algorithms.

Since these TOF arcs do not perfectly intersect in practice (e.g., due to noise), an estimation method has to be used; often, this method has to simultaneously solve a data association problem (Too *et al.*, 2017). When the locations of reflectors are known, multipath can be leveraged for better accuracy. Reflections can be interpreted as producing ‘virtual receivers’ as illustrated in Fig. 4, which corresponds to the structure of our watertank setting.

We make the simplification that reflective structures are (piecewise) approximated as planar boundaries (Deane, 1994), characterized by the range ρ , azimuth θ , and elevation ϕ of their normal vector relative to the origin. Thus, boundary i is parametrized as the vector $\boldsymbol{\eta}_i = [\rho_i, \theta_i, \phi_i]^T$. We denote the Cartesian coordinates of the emitter and receiver positions as $\mathbf{p}_e \in \mathbb{R}^3$ and $\mathbf{p}_r \in \mathbb{R}^3$, respectively. The LOS TOF, τ_0 , is calculated as

$$\tau_0 = \frac{\|\mathbf{p}_r - \mathbf{p}_e\|_2}{v_s}, \tag{1}$$

where v_s is the constant speed of sound, which can be a valid approximation for well-mixed shallow-water setting at short ranges and high acoustic frequencies (Chitre, 2007;

Too *et al.*, 2017). For boundary i , we obtain the virtual receiver location \mathbf{p}_i by finding the corresponding reflection of \mathbf{p}_r . The TOF to the receiver from the NLOS arrival of boundary i is equal to the TOF from \mathbf{p}_e to the corresponding \mathbf{p}_i ; thus

$$\tau_i = \frac{\|\mathbf{p}_i - \mathbf{p}_e\|_2}{v_s}. \tag{2}$$

We denote a single copy of the periodically emitted signal as $s(t)$, and we merge the effects of attenuation and reflection as the equivalent attenuation coefficient α_i for each path. The received signal $r(t)$ is then

$$r(t) = \sum_{i=0}^N \alpha_i s(t - \tau_i(\mathbf{p}_e, \boldsymbol{\eta}_i)) + n(t), \tag{3}$$

where $n(t)$ is an unknown random noise realization, and N is the (known) number of reflective surfaces. Note that α_i can be zero, as when an occluder blocks a particular arrival, so Eq. (3) incorporates occlusion. We denote each receiver as $\mathbf{p}_{r,j}$, each of which has TOFs $\tau_{i,j}$ from their own virtual receivers, and received signals $r_j(t)$ [defined as in Eq. (3)].

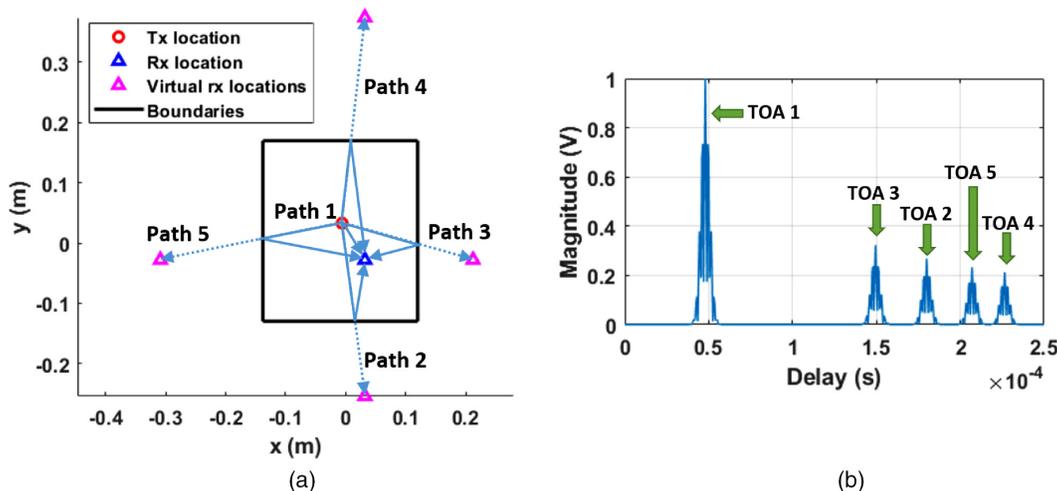


FIG. 4. (Color online) Illustration of virtual receivers produced by reflective planar boundaries in an occluder-free environment (a), and the corresponding noise-free simulated received signal for this scenario (b). A NLOS TOF to the real receiver is equivalent to the LOS TOF to the virtual receiver.

We only model the single-bounce reflections, which are typically the strongest multipath components. In practice, not only will there be reverberation, but the N that is assumed known may be inaccurate. A proposed localization method must be robust to the resulting discrepancies.

We model $s(t)$ as pulsed and periodic, which limits our estimation scenarios to such man-made sources. This model allows us to record $s(t)$ and to infer some key parameters: the approximate (truncated) channel impulse response length T_c (Sun and Wang, 2018), and the time period T_p between emissions. We require $s(t)$ to meet the following constraints in order to have accurate estimates $\{\hat{\tau}_{i,j}\}$ of the $\{\tau_{i,j}\}$. First, we define T_s as the temporal duration of $s(t)$, so that $s(t)$ is strictly zero outside $0 \leq t \leq T_s$. Next, we define T_A to be the full width at half maximum (FWHM) of the autocorrelation function of $s(t)$, which is the time resolution of the matched-filtered emitted signal. Finally, we define T_m as the minimum time difference between the main multipath arrivals. The PEEL operation modes are then as follows:

- If $T_s < T_m$, we have a pulse that is shorter than the multipath separation, as is the case with the Gaussian pulse in our experiments. Here, we do not need to know $s(t)$ for PEEL; T_s is short enough that the multipath arrivals are roughly separable.
- If $T_s > T_m$ and $T_A < T_m$, then after matched-filtering, the multipath peaks are roughly separable. Now, we do in fact need to know $s(t)$ for this matched-filtering step, so that we are able to identify the distinct arrivals and thereby solve for the environment.

We also need $T_p > T_c$ to avoid different periodic emissions being observed within the same $r(t)$.¹ With these introductory developments, we are ready to present an architecture for reliable localization in the presence of such errors.

III. END-TO-END LOCALIZATION AND STRUCTURE LEARNING

We now present our proposed joint localization and structure learning method, PEEL. First, we use TDOA and TOA localization for initialization; then obtain an initial boundary localization estimate; and finally track a moving emitter as we simultaneously learn the environment, as summarized in Fig. 5.

A. Stage 1: Coarse localization of the emitter

We first describe a standard TDOA localization algorithm (which does not require knowledge of $\{\eta_i\}$) for synchronization with the emitter (Korhonen, 2008), then a TOA algorithm that leverages multipath from known boundaries $\{\eta_i\}$ after synchronization (Ribeiro *et al.*, 2010; Brutti *et al.*, 2010). Both algorithms use a grid-search for the emitter to obtain $\hat{\mathbf{p}}_e$, where a set \mathcal{P} of grid points \mathbf{p}_{cand} is defined over the target area.² Each $\mathbf{p}_{\text{cand}} \in \mathcal{P}$ corresponds to a potential emitter location, with an associated set of $\{\tau_{i,j,\text{cand}}\}$ to the receivers $\mathbf{p}_{r,j}$ and the corresponding virtual receivers $\mathbf{p}_{i,j}$.

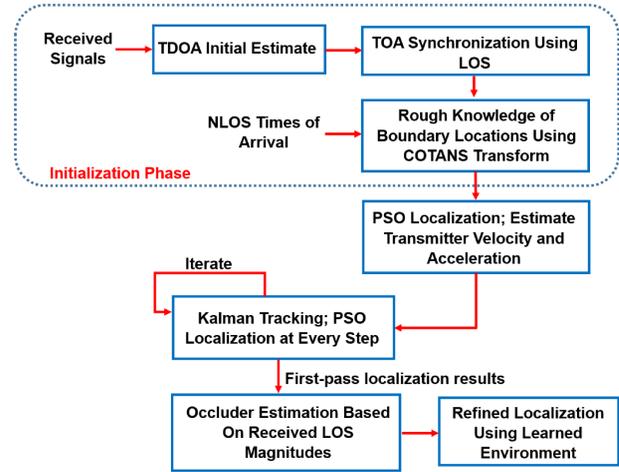


FIG. 5. (Color online) Conceptual flow chart for the proposed PEEL architecture.

Thus, each $\tau_{i,j,\text{cand}}$ is a function of $\mathbf{p}_{r,j}$ and of η_i . The TDOA and TOA stages are used to initialize more refined techniques.

TDOA does not require a common clock with the emitter and is suitable to establish synchronization (Su *et al.*, 2017; Qiao *et al.*, 2013). Denoting the calculated cross correlation between the signals at receivers j_1 and j_2 as $R_{j_1,j_2}(\tau)$, the TDOA metric that we use is

$$C_{\text{TDOA}}(\tau_{\text{cand}}(\mathbf{p}_{\text{cand}})) \triangleq \sum_{j_2 \neq j_1} \sum_{j_1=1}^M |R_{j_1,j_2}(\tau_{j_1,\text{cand}}(\mathbf{p}_{r,j_1}) - \tau_{j_2,\text{cand}}(\mathbf{p}_{r,j_2}))|. \quad (4)$$

Using the generalized cross correlation with phase transform (GCC-PHAT) method for better performance (Knapp and Carter, 1976), we estimate the emitter location $\hat{\mathbf{p}}_{e,\text{TDOA}}$ as

$$\hat{\mathbf{p}}_{e,\text{TDOA}} = \underset{\mathbf{p}_{\text{cand}} \in \mathcal{P}}{\text{argmax}} C_{\text{TDOA}}(\tau_{\text{cand}}(\mathbf{p}_{\text{cand}})). \quad (5)$$

The average difference between LOS peak delays and the expected TOFs from TDOA is then used to establish synchronization.³ One can then use more accurate TOA localization methodologies. We define the additive TOA metric $C_{\text{TOA}}(\tau_{\text{cand}}(\mathbf{p}_{\text{cand}}))$ as

$$C_{\text{TOA}}(\tau_{\text{cand}}(\mathbf{p}_{\text{cand}})) \triangleq \sum_{j=1}^M \sum_{i=0}^N |r_j(t - \tau_{i,j,\text{cand}}(\mathbf{p}_{r,j}, \eta_i))|, \quad (6)$$

where the time indices $(t - \tau_{i,j,\text{cand}}(\mathbf{p}_{r,j}, \eta_i))$ are rounded to the nearest sample in a discrete-time implementation. The estimate $\hat{\mathbf{p}}_{e,\text{TOA}}$ is then obtained as

$$\hat{\mathbf{p}}_{e,\text{TOA}} = \underset{\mathbf{p}_{\text{cand}} \in \mathcal{P}}{\text{argmax}} C_{\text{TOA}}(\tau_{\text{cand}}(\mathbf{p}_{\text{cand}})). \quad (7)$$

The heuristic metric in Eq. (6) is maximized when each $\tau_{i,j}$ corresponds to a peak location; such locations are the intersections of TOF arcs, as seen in Fig. 6(a). When the

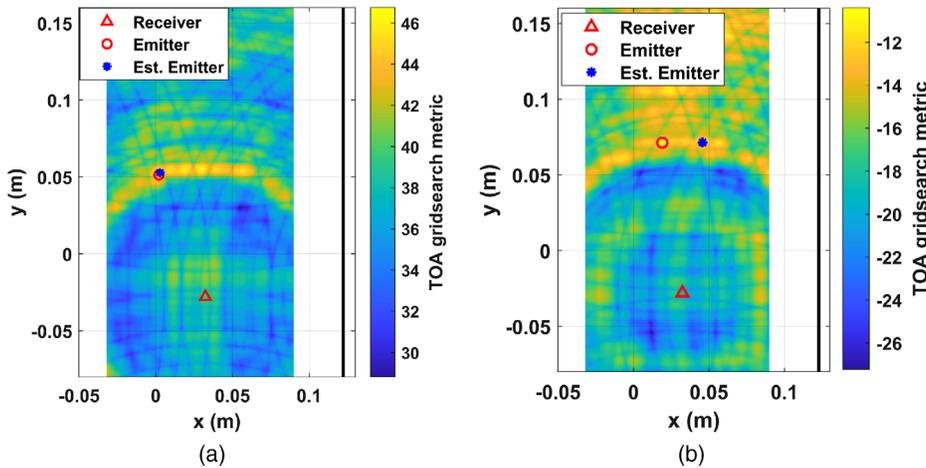


FIG. 6. (Color online) A TOA grid-search that localizes the emitter on real data (a), and a grid-search that fails at a different location due to the alternative arrival arc crossings (b), at ~ 36 dB SNR.

$\{\eta_i\}$ are known, and therefore the virtual receivers are labeled correctly, $\hat{\mathbf{p}}_e$ is obtained accurately.

TDOA and TOA localization are more robust if the search space \mathcal{P} can be limited to the vicinity of the correct location, \mathbf{p}_e . Consider the results in Fig. 6(a); although \mathbf{p}_e is estimated accurately, there are many local maxima in the search space. These can lead to large errors as in Fig. 6(b). Confining \mathcal{P} to the vicinity of \mathbf{p}_e avoids such ambiguities. We summarize our TDOA and TOA localization algorithms in Algorithms 1 and 2.

One advantage of Algorithm 2 is that it is a heuristic weighting of the arrivals. A localization method that is solely based on $\{\tau_{i,j}\}$ will assign equal importance to each arrival so that there will be no automatic compensation for faulty TOF estimates. A TOF estimate is likelier to have a high error if the corresponding arrival is weak. In such cases, its contribution to C_{TOA} in Algorithm 2 will be smaller, and its impact on $\hat{\mathbf{p}}_{e,\text{TOA}}$ in Eq. (7) will be reduced.

ALGORITHM 1. TDOA localization for initialization.

Require: Unoccluded LOS arrivals at each receiver.

Input: \mathcal{P} , $\{r_j(t)\}_{j=1}^M$, $\{\mathbf{p}_{r,j}\}_{j=1}^M$

for $\mathbf{p}_{\text{cand}} \in \mathcal{P}$ **do**

 Calculate C_{TDOA} [Eq. (4)]

end

Get $\hat{\mathbf{p}}_{e,\text{TDOA}}$ [Eq. (5)]

Result: $\hat{\mathbf{p}}_{e,\text{TDOA}}$

ALGORITHM 2. TOA localization after initialization.

Require: Time synchronization after applying Algorithm 1; number of boundaries N is known.⁴

Input: \mathcal{P} , $\{r_j(t)\}_{j=1}^M$, $\{\mathbf{p}_{r,j}\}_{j=1}^M$, $\{\eta_i\}_{i=1}^N$

for $\mathbf{p}_{\text{cand}} \in \mathcal{P}$ **do**

 Calculate C_{TOA} [Eq. (6)]

end

Get $\hat{\mathbf{p}}_{e,\text{TOA}}$ [Eq. (7)]

Result: $\hat{\mathbf{p}}_{e,\text{TOA}}$

Algorithm 2 also allows us to compare different hypothesized environments. Given candidate boundaries $\{\eta_i\}_{\text{cand}}$, we can run Algorithm 2 for each $\{\eta_i\}_{\text{cand}}$, and pick the maximum metric solution as our joint estimate of $\hat{\mathbf{p}}_{e,\text{TOA}}$ and $\{\eta_i\}$. C_{TOA} is therefore a heuristic measure of the estimate's goodness. Recall that the $\{\eta_i\}$ cannot generally be assumed known; thus, we next propose a boundary localization algorithm to produce a suitable set of $\{\eta_i\}_{\text{cand}}$.

B. Stage 2: Preliminary boundary localization

After TDOA and TOA localization, we have an initial estimate of the emitter's position and coarse synchronization. To jointly estimate the emitter and the environment in a computationally feasible manner, we need a good initial estimate of reflective boundary positions, which ultimately yields a tight search space for $\{\eta_i\}$. We assume that a set of unlabeled $\{\hat{\tau}_{i,j}\}$ has been estimated (Demirli and Saniie, 2001), though some of these estimated arrival times might have large errors. We now develop a Hough transform-inspired method to estimate $\{\eta_i\}$ which is robust to errors in $\{\hat{\tau}_{i,j}\}$, building on (Naseri and Koivunen, 2016).

As some background context, Euclidean distance matrices (EDM) (Dokmanic et al., 2015) and convex optimization methods (Naseri et al., 2014) have been used to solve for virtual emitter positions and thereby obtain boundary positions. These methods can suffer from large errors if an arrival $\hat{\tau}_{i,j}$ has been labeled incorrectly, and their performances are initialization-sensitive. An alternative approach relies on the fact that in two dimensions, these $\hat{\tau}_{i,j}$ each define an ellipse whose foci are \mathbf{p}_e and \mathbf{p}_r , as in Fig. 7(a); in three dimensions, they define spheroids. In Naseri and Koivunen (2016), boundaries are observed to be common tangents to such ellipses; by fitting tangent planes to these ellipses, we can estimate the boundaries while *avoiding* the error-prone initialization and labeling procedure in multipath environments, such as in Fig. 7(b). It is convenient to define a given η_i by the spherical coordinates of a vector from the origin (arbitrarily selected) that is orthogonal to the boundary plane. Thus, each boundary is specified by a range ρ , azimuth θ , and elevation ϕ , where these local spherical

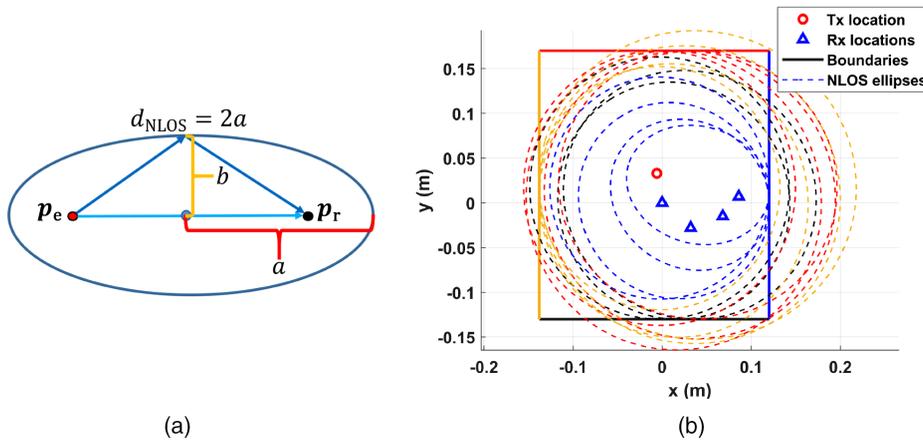


FIG. 7. (Color online) The NLOS arrivals define ellipses/spheroids of equidistance (a); we illustrate the ellipses defined by $\{\tau_{i,j}\}$ for our experimental setting (depicted in 2D for convenience). Each boundary is a common tangent to a single ellipse due to each receiver, highlighted here by matching colors.

coordinates correspond to a standard east-north-up Cartesian coordinate system.

The transform for boundary localization in Naseri and Koivunen (2016), which for convenience we refer to as common tangents to spheroids (COTANS), is derived as follows. A boundary defined by ρ , θ , and ϕ can be conceptualized as a point (ρ, θ, ϕ) in a COTANS transform domain; working out the (ρ, θ, ϕ) expression of a plane is to take the plane's COTANS transform (Borrmann et al., 2011). When the space $\rho \times \theta \times \phi$ is discretized as a 3-mode tensor, each hypothetical tangent has a finite set of (ρ, θ, ϕ) on which it may lie. Incrementing this “accumulator” tensor over every potential (ρ, θ, ϕ) entry yields a COTANS-domain image [as in Fig. 9(a)], with maxima at the true boundaries in the noiseless case.

Here, we extend the method of Naseri and Koivunen (2016) to three dimensions, and we further derive a direct solution for the tangents in the COTANS domain, which precludes the need for randomly sampling points (ρ, θ, ϕ) on the surface of an ellipse/spheroid. We define a grid over θ and ϕ , and obtain the COTANS transform for the tangents at each θ and ϕ . In two dimensions, for a given θ , we obtain the COTANS transform's ρ in Appendix A 1 as

$$\rho(\theta) = \sqrt{a^2 \cos^2 \theta + b^2 \sin^2 \theta}, \quad (8)$$

and in three dimensions, the result for a given θ and ϕ is calculated in Appendix A 2 as

$$\rho(\theta, \phi) = \sqrt{c^2 \cos^2 \phi + a^2 \sin^2 \phi}, \quad (9)$$

where a , b , and c are standard axes of spheroids calculated from the $\{\hat{\tau}_{i,j}\}$. This result is obtained for an origin-centered spheroid. To move a point cloud of (ρ, θ, ϕ) centered on the origin to the true emitter and receiver positions, we rotate the points to match the true spheroid's orientation (adjusting their θ and ϕ to be some θ_{rot} and ϕ_{rot}), and then translate the points to the true spheroid's position (yielding a final ρ_{COTANS} , θ_{COTANS} , and ϕ_{COTANS}). These COTANS-domain operations are illustrated in Fig. 8.

In practice, we first generate a set of points $\{(\rho, \theta, \phi)\}$ by defining an azimuth-elevation grid, and obtaining the COTANS transforms for each of the corresponding tangents. The resulting points $\{(\rho, \theta, \phi)\}_{\text{COTANS}}$ are then rounded to a desired accuracy.⁵ We define a tensor over ρ , θ , and ϕ with this resolution, for each rounded $\{(\rho, \theta, \phi)\}_{\text{COTANS}}$ point, we increment the corresponding cell in the tensor by 1. We sum the tensors for each NLOS arrival at each receiver, to obtain a final tensor of superimposed COTANS curves such as in Fig. 9(a). Each ellipse contributes a curve; observe that the curves do not intersect at a single point because of errors in $\{\hat{\tau}_{i,j}\}$ and in the measured \mathbf{p}_e and \mathbf{p}_r , so we have multiple maxima. We therefore use a smoothing filter to replace the tensor values with local averages.⁶ Then, we extract the locations of as many maxima as there are boundaries as in Fig. 9(b), where we set the neighborhood of each maximum

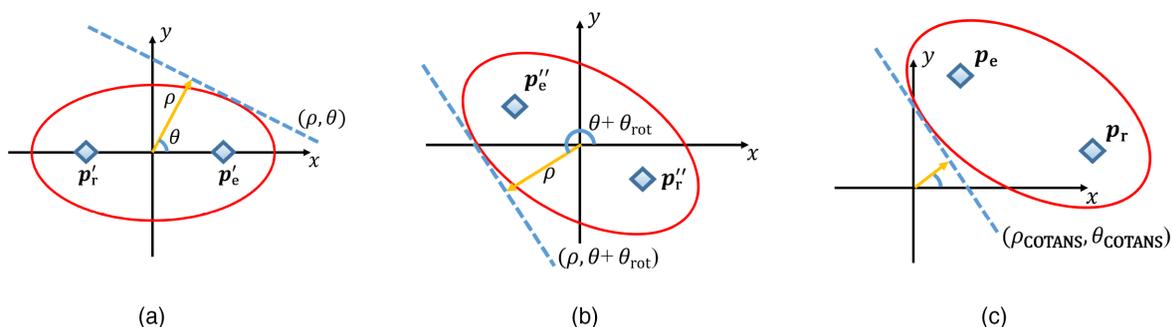


FIG. 8. (Color online) Illustration of the steps to obtain the COTANS transform of a particular tangent line (depicted in 2D). Description of one tangent to a standard ellipse (a); rotation of this origin-centered ellipse and its tangent (b); and the translation of this ellipse to its real position (c).

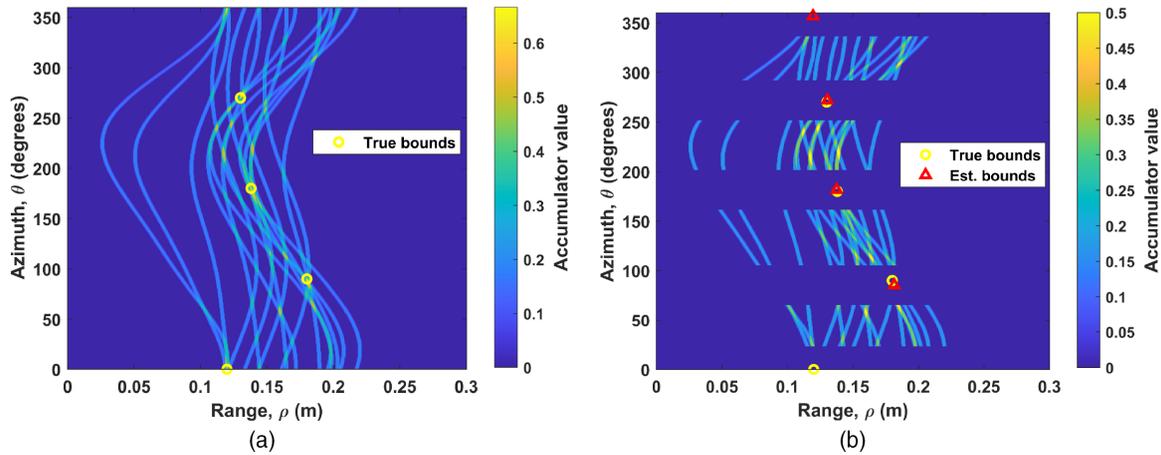


FIG. 9. (Color online) The COTANS accumulator for inconsistent $\{\hat{\tau}_{i,j}\}$ (a), and its boundary estimates (b). The image is periodic in azimuth, so the maximum near 360° is close to the correct location at 0° .

ALGORITHM 3. COTANS transform-based boundary localization.

Require: Time synchronization from Algorithm 2; number of boundaries N is known.

Input: $\{\hat{\tau}_{i,j}\}, \{\mathbf{p}_{r,j}\}_{j=1}^M, \hat{\mathbf{p}}_{e,\text{TDOA}}$

for $\hat{\tau}_{i,j} \in \{\hat{\tau}_{i,j}\}$ **do**

Calculate COTANS transform for spheroid, add to accumulator (Appendix A)

end

Result: Apply smoothing to COTANS accumulator, find its N maxima as $\{\hat{\boldsymbol{\eta}}_i\}$.

in the tensor to zero, to avoid picking the same boundary multiple times.⁷ This method is summarized in Algorithm 3.

Modeling the environment as consisting of planar boundaries makes environment learning more computationally tractable. Note that the virtual receiver model is not limited to planar boundaries: the planar boundary estimates can be tangents to a non-planar boundary. In this case, however, the boundary estimates will be inaccurate for localization purposes when the emitter moves to a different position. Therefore, rather than using the initial COTANS boundary estimates for the entire emitter trajectory, we need evolving boundary estimates over time, utilizing the global optimization method which we now discuss.

C. Stage 3: Joint boundary estimation and emitter localization

At this stage, we have initial estimates of both the emitter position and the boundaries, allowing for simultaneous optimization over \mathbf{p}_e and $\{\hat{\boldsymbol{\eta}}_i\}$. We assume candidate boundaries $\{\boldsymbol{\eta}_i\}_{\text{cand}}$ that are in the vicinity of the $\{\hat{\boldsymbol{\eta}}_i\}$, and then carry out TOA localization (Algorithm 2) for each $\{\boldsymbol{\eta}_i\}_{\text{cand}}$. These $\{\boldsymbol{\eta}_i\}_{\text{cand}}$ are drawn from a search space $\rho \in [\rho_i - \rho_0, \rho_i + \rho_0]$, $\theta \in [\theta_i - \theta_0, \theta_i + \theta_0]$, $\phi \in [\phi_i - \phi_0, \phi_i + \phi_0]$, where $\rho_i, \theta_i,$ and ϕ_i are estimated from Algorithm 3, and $\rho_0, \theta_0,$ and ϕ_0 are error margins obtained as the typical maximum parameter estimation errors from COTANS, for a given setting. The $\{\hat{\boldsymbol{\eta}}_i\}$ associated with the highest maximum metric

C_{TOA} [Eq. (6)] yields the most likely \mathbf{p}_e and $\{\boldsymbol{\eta}_i\}$. Interlacing this algorithm with tracking over the emitter positions, we discover and map the occluders in the environment, and refine our estimates for joint localization and environment learning.

In this simultaneous optimization problem with measurement errors, it is advantageous to maintain multiple evolving estimates of the emitter and boundary locations, for which we use particle swarm optimization (PSO) (Clerc, 2010). In this global optimization algorithm, each $\{\boldsymbol{\eta}_i\}_{\text{cand}}$ is termed a ‘‘particle.’’ Each particle is initialized at a random location and given an initial ‘‘velocity’’ in the search space. Algorithm 2 is carried out for each particle, and the maximum C_{TOA} and corresponding $\hat{\mathbf{p}}_{e,\text{cand}}$ is recorded. Particle velocities are modified at each PSO iteration towards the global maximum metric particle (with the maximum C_{TOA}), and also to the maximum metric particle in a random ‘‘neighborhood’’ of particles. Thus, particles explore the region around the current best approximation.

PSO was used in PEEL because its metaparameters do not need to be adapted to a specific physical setting and because it can be run without re-initialization during tracking. There are alternative global optimization algorithms that can also be used in PEEL, however, and it would be advantageous in an application setting to fine-tune the metaparameters of PSO or another algorithm for that particular environment.

In order to be able to accommodate for occlusion, we rely on a tracking procedure over the periodic transmissions of the moving emitter. Kalman tracking with a constant-acceleration model, for example, can be applied. The $\hat{\mathbf{p}}_e$ obtained using PSO is used as the measurement vector to update the tracking. Critically, the grid-search is confined to a sufficiently small area \mathcal{P} around the resulting prediction. This way, even if the emitter is temporarily occluded and the underlying C_{TOA} metric is erroneous, $\hat{\mathbf{p}}_e$ will still be close to the actual \mathbf{p}_e . Thus, if the emitter enters an area that is unfavorable for localization, PEEL can potentially mitigate the resulting estimation errors. Our algorithm for PSO localization and boundary estimation is given in Fig. 10, and our overall approach is summarized in Algorithm 4.

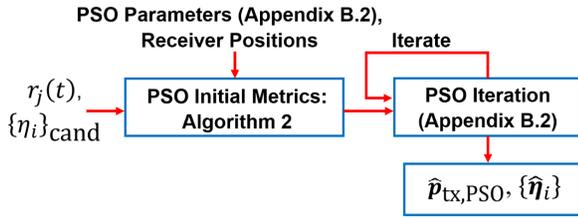


FIG. 10. (Color online) Flowchart for simultaneous boundary and emitter localization with PSO.

ALGORITHM 4. Joint PSO localization and Kalman tracking.

Input: $\{r_j(t)\}_{j=1}^M, \{\mathbf{p}_{r,j}\}_{j=1}^M, \{\eta_i\}_{\text{cand}}$
Require: Moving emitter.

for each emitter position in trajectory do
 for pre-specified number of PSO iterations do
 Perform PSO, get $\hat{\mathbf{p}}_{\text{e,PSO}}, \{\hat{\eta}_i\}$ (Fig. 10, Appendix B)
 end for
 Use $\hat{\mathbf{p}}_{\text{e,PSO}}$ to update Kalman tracking, get new \mathcal{P} for next iteration.
 end for

Result: $\hat{\mathbf{p}}_{\text{e,PSO}}, \{\hat{\eta}_i\}$

From the results of Algorithm 4, we can determine whether each LOS path to receivers has been attenuated beyond expectation. Occlusion can be determined by assuming a $1/\hat{R}$ attenuation factor in the received signal magnitude, where \hat{R} is a given estimated LOS range. Unoccluded arrivals allow estimation of the emitted power, by scaling the corresponding LOS magnitudes by \hat{R} . If the received LOS magnitude is substantially below the magnitude predicted by geometric spreading, we conclude that an occluder is present.⁸ Volumetric intersection methods can then be used to map the occluders (Srinivasan *et al.*, 1990).

After mapping the occluders, we re-run Algorithm 4. We now know which LOS arrivals are blocked by the occluder, and we change the additive metric C_{TOA} by leaving blocked paths out of the sum in Eq. (6). Our flexible cost function thereby allows us to compensate for model mismatch due to occlusion, and yields increasingly improved performance.

IV. EXPERIMENTAL RESULTS

We now evaluate our PEEL method for its passive localization and environment learning capabilities, seeking to demonstrate and understand its performance and robustness. For this purpose, we use the convenient setting of the watertank depicted in Fig. 2. We first test PEEL on real data obtained from the watertank; then add synthetic noise to this data to analyze PEEL’s performance compared to other localization methods and finally conduct trials in a simulation setting with a wider range of localization scenarios.

In our experiments, we emit a high-frequency Gaussian pulse

$$s(t) = e^{-t^2/f_0^2} \cos(2\pi f_0 t), \tag{10}$$

where the center frequency f_0 was set to 280 kHz, corresponding to a wavelength of 5 mm. The boundaries used for

localization are the water surface, whose position is known since the receiver depths are known; and the bottom and three plastic sides of the tank, whose ranges to the origin and azimuths are respectively known to within ± 5 mm and $\pm 6^\circ$, as obtained using COTANS. One side-boundary is deliberately left unaccounted for, which presents the difficulty of having an unmodeled set of arrivals in the data so that we have $N=5$. We use $M=4$ receivers. The speed of sound through water in the tank was measured as 1485 m/s.

For our PEEL implementation, we use 128 particles and 5000 PSO iterations for each emitter position. At each PSO iteration, we use a grid \mathcal{P} of 0.5 mm resolution and sides of 6 mm centered on the expected emitter location. We perform ten experiments for each localization scenario and obtain the root mean square (rms) localization error. The rms error for the first pass of PEEL is 3.9 mm, as seen in Fig. 11(a). To put this into perspective, the average distance of the emitter from the receivers for this trajectory is 93.5 mm, and the rms error in this scenario is 4.2% of the mean distance to the emitter. The final boundary estimates in Fig. 11(a) are obtained using the full estimated trajectory, which leads to a more accurate solution compared to the results obtained from using just one emitter position. Running PEEL’s second pass, we get the result in Fig. 11(b), where the rms error is reduced to 2.7 mm (2.9% of the mean distance). The occluder is reconstructed with sufficient accuracy to improve our localization result. To understand how much the unknown occluder’s presence affects the result, we experimented with the same scenario but with no occluder, obtaining an rms error of 1.5 mm (1.6% of the mean distance), which is within our margin of measurement accuracy. In these trials, occlusion and interference due to reverberation are the key error factors; the SNR of unoccluded LOS arrivals is > 30 dB.

If synchronization and an initial boundary estimate are given to us, single-receiver localization using Algorithm 4 becomes feasible. In an experiment with a single emitter at the origin and no occluder, we get 2.1 mm of rms error (2.2% of the mean distance). This strong performance is surprisingly only slightly worse than when we use four receivers.

In order to analyze the performances of the localization algorithms that we have examined (TOA localization with known boundaries, and PSO), we carried out localization trials with synthetic Gaussian noise added to experimental signals. For the first position in the trajectory, where there is LOS to every receiver, the same noise power in units of dBm was added to each recorded signal. In Fig. 12(a), we compare PSO’s performance for different numbers of PSO iterations, vs that of TOA localization with known boundaries. Because PSO optimizes for the boundary locations as well, its high runtime led to us conduct 150 trials for each noise level, while 5000 trials were carried out for TOA. There is a performance gap between PSO and the “ground truth” TOA result, which is smaller for higher noise levels. Longer runtimes improve PSO’s performance and help close this gap.

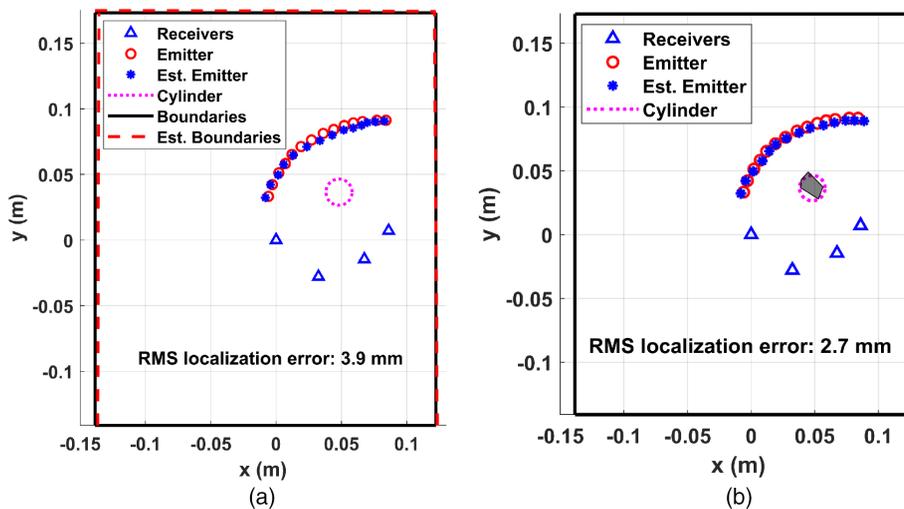


FIG. 11. (Color online) First-pass localization performance, with an occluding cylinder present in the environment (a), and the improved second-pass result, using the modified metric obtained by taking the occluder into account (b). The boundary that is unaccounted-for to introduce model mismatch is the bottom boundary in these top-view figures.

We next evaluate the performance of environment learning by adding synthetic Gaussian noise to the signals measured at the last position in the trajectory (where two LOS arrivals are occluded) and obtaining the rms range error from the origin to all the boundaries, averaged over trials and the number of boundaries. We conducted 100 such trials for each noise level and for three scenarios: in the absence of the occluding cylinder as measured in a separate set of experiments, with the occluding cylinder present but not accounted for, and with the occluded arrivals removed in the revised metric (modeling the PEEL second pass). The results are given in Fig. 12(b), where we observe a gap in performance caused by the occlusion, beyond the performance deterioration caused by a lower effective SNR. Boundary estimation is improved by compensating for the occluder, though this improvement is less dramatic than the performance gains in emitter localization. Adding a noise power greater than 20 dBm causes much larger errors in the localization and boundary estimation tasks in this setting, as the noise peaks have much greater magnitude than the NLOS arrivals.

To test PEEL’s performance on a wider range of experimental geometries, we generated synthetic data that modeled the watertank setting with randomized receiver locations. We implemented a simulation setting as per Deane (1994) that modeled a reverberant environment with multibounce arrivals, producing synthetic received signals with the mirror image method. In these trials, the emitter’s trajectory, the boundary locations, the occluder, and the position of the receiver at the origin were kept constant, and PEEL was run for 10 different deployments of the other three receivers. These receivers were distributed uniformly over the region bounded by the positions of receivers in the real experimental setting. We used the same Gaussian pulse that was used in the watertank experiment, though the only distortion that these pulses underwent in simulation was attenuation. Ultimately, we obtained an rms error of 2.6 mm over ten trials, which was close to the watertank results.

Overall, we observed reliable performance across all of the trials for this particular demonstration. Once the unknown occluder is compensated-for, we only observe several millimeters of error, which is reasonable since there are

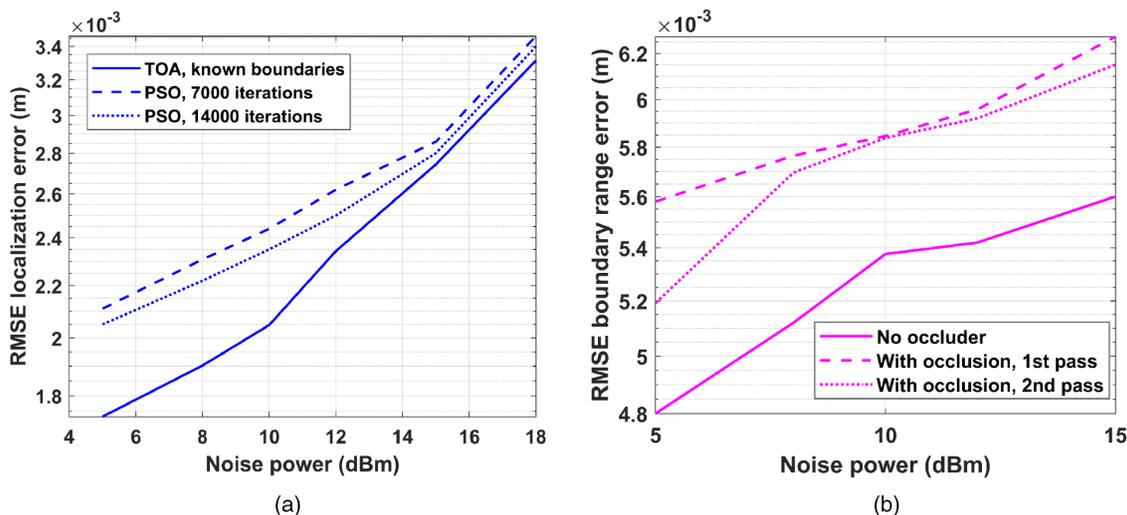


FIG. 12. (Color online) Performance of PSO localization versus TOA with known boundaries, highlighting better accuracy with more PSO iterations (a); and rms boundary localization error performances for experimental results with synthetic added noise (b).

TABLE II. Localization experiment results.

Scenario	rms Error (mm)
Four receivers, no occluder	1.5
Single receiver, no occluder	2.1
Four receivers, occluder, first-pass	3.9
Four receivers, occluder, final result	2.7
Simulation, random receivers	2.6

already millimeter-scale errors in the assumed receiver positions. These results are summarized in Table II.

V. CONCLUDING REMARKS

In this paper, we developed a method for localization and environment learning, which was tested in a three-dimensional (3D) reverberant underwater environment with occluding objects. Our PEEL method performed iteratively refined localization, with each solution yielding results that served as assumed knowledge for the following more sophisticated algorithm stages. This multi-stage approach handled a complex problem where successively establishing synchronization, boundary estimation, emitter tracking, and occluder estimation, made solving the overall task computationally feasible. Applying these stages separately is suboptimal as compared to a hypothetical joint optimization over all of the aspects of the problem. However, PEEL breaks down the localization task into subproblems with well-posed solutions. We ultimately achieved accurate simulation and experimental results in end-to-end localization.

The trajectory to which PEEL was applied resulted in a difficult localization problem which nevertheless proved to be solvable. There are many adversarial scenarios, however, that can defeat our method. The emitter may move behind an occluder and come to a halt, or it may take a sharp turn which throws off the tracking algorithm. One way to cope with such scenarios would be to run several instances of PEEL in parallel. A discrepancy between their results would indicate catastrophic errors. This parallelization strategy for PEEL is also important for choosing the number of boundaries, N . If one is not confident about the assumption on N , then running PEEL in parallel with different values of N becomes advantageous. Choosing the best N is an important open problem for PEEL.

Note that we do not require a complete deduction of the occluders for improved localization performance. In this passive setting, a complete occluder mapping requires the emitter to move in an arc around them, as is the case in tomography. It is not guaranteed that occluders can be mapped since the emitter’s path is not under our control. As long as occlusion is detected, potential errors from summing over nonexistent arrivals are avoided.

When deploying PEEL in an open water setting, there are improvements that must be made to cope with this more difficult environment. To use as general a model as possible, we have used the simplest versions of each algorithm, such as GCC-PHAT for TDOA, basic PSO for optimization, and a

constant acceleration model for Kalman tracking. Each of these algorithms can be enhanced if more domain knowledge is available. Applying an algorithm that tracks the multipath arrivals (Fleury *et al.*, 1999; Salmi *et al.*, 2008), for example, can help to further refine the boundary estimates. The underwater environment’s bathymetric profile or the presence of occluding objects may already be known, and must then be taken into account. PEEL can also be applied to tasks such as indoor (acoustic or RF) localization, which features easier channels, but cluttered environments with different challenges.

PEEL currently incorporates global optimization over the emitter position and the locations of planar boundaries. A more sophisticated environmental model would allow it to work robustly with the complex non-planar boundaries in the real-life shallow-water environment. In the literature, global optimization with simulated annealing and genetic algorithms have been applied to MFP (Hermand, 1999) for geoacoustic inversion, using known emitters. Optimization over not just the locations of their boundaries but their physical properties, while more computationally expensive, could lead to improved results in field deployments.

There are broader questions in TOA localization that pertain to PEEL. Currently, PEEL exploits the emission of deterministic, pulsed signals. In practice, other types of emissions may arise, as with communications sources or vessel signatures. It would be challenging to establish time synchronization with such signals. Although we have experimented with performance in increasing noise, it would also be of interest to obtain a full bandwidth/noise characterization over different environments and signals. Finally, COTANS boundary estimation may be amenable to a machine learning method, to automate the smoothing and peak extraction tasks. COTANS transform results are two-dimensional (2D) or 3D images, so we can train a neural network on such data for boundary estimation in future experiments.

ACKNOWLEDGMENTS

This work was supported, in part, by ONR under Grant Nos. N00014-19-1-2661, N00014-19-1-2662, and N00014-19-1-2665, and NSF under Grant No. CCF-1816209. The authors also thank Dr. Jim Preisig of JP Analytics LLC for valuable discussions and feedback on the work.

APPENDIX A: DERIVATION OF THE BOUNDARY LOCALIZATION APPROACH

1. COTANS transform for tangent lines to ellipses in 2D

We present the derivation of the COTANS transform of tangent lines to an ellipse in 2D. First, we derive the tangent line to a standard ellipse centered around the origin, with a closed-form description in terms of the normal vector to the tangent. We then consider this ellipse to be rotated and translated to its actual position in 2D and modify the normal-vector description of the corresponding rotated and translated tangent. The final parametrization of the tangent in terms of its normal vector is $(\rho_{\text{COTANS}}, \theta_{\text{COTANS}})$; this is a

point in $\{(\rho, \theta)\}$ -space, where ρ is the length of the vector and θ is its direction as measured with respect to the x -axis. This description is by definition the COTANS transform of the tangent line.

First, we obtain a description of tangent lines to standard ellipses [as in Fig. 7(a)]. The equation for a standard ellipse centered on the origin is $(x^2/a^2) + (y^2/b^2) = 1$, where a and b are its major and minor axes, respectively. Consider a point (x_0, y_0) on the ellipse, and note that this point can be parametrized as $(a \cos \alpha, b \sin \alpha)$, where the angle α is an intermediate parametrization. Differentiating the ellipse's equation, the slope of a tangent line is

$$\frac{2x}{a^2} + \frac{2y}{b^2} \frac{dy}{dx} = 0 \Rightarrow \frac{dy}{dx} = -\frac{b^2 x}{a^2 y}. \tag{A1}$$

From the fact that (x_0, y_0) satisfies both $(x_0^2/a^2) + (y_0^2/b^2) = 1$ and the equation of the tangent's slope, we obtain the equation of the tangent line at (x_0, y_0) as $(xx_0/a^2) + (yy_0/b^2) = 1$.

Now, we obtain the COTANS transform of this tangent line, by parametrizing it in terms of its Euclidean distance ρ from the origin and the azimuth θ of the normal line to the tangent going through the origin. Thus, a tangent to the ellipse at $(a \cos \alpha, b \sin \alpha)$ will be expressed as a point $(\rho_\alpha, \theta_\alpha)$ in the COTANS domain. One derivation of this COTANS transform result is as follows. Our method will obtain $\rho(\theta)$, the distance to the origin of a given tangent plane as a function of azimuth.

The distance ρ_α of the tangent line to the origin is the solution to

$$\begin{aligned} \min \quad & \|\mathbf{x}\|_2 \\ \text{s.t.} \quad & \mathbf{w}_\alpha^T \mathbf{x} = 1, \end{aligned} \tag{A2}$$

where $\mathbf{x} = [x \ y]^T$ is the point of intersection of the tangent line and its normal through the origin, and $\mathbf{w}_\alpha^T = [x_0/a^2 \ y_0/b^2] = [\cos \alpha/a \ \sin \alpha/b]$. ρ_α can be found by using Lagrange multipliers technique, where

$$\begin{aligned} \mathcal{L}(\mathbf{x}, \lambda) &= \frac{1}{2} \mathbf{x}^T \mathbf{x} - \lambda (\mathbf{w}_\alpha^T \mathbf{x} - 1), \\ \nabla_{\mathbf{x}} \mathcal{L}(\mathbf{x}, \lambda) = \mathbf{0} &\Rightarrow \mathbf{x} - \lambda \mathbf{w}_\alpha = \mathbf{0} \Rightarrow \mathbf{x} = \lambda \mathbf{w}_\alpha, \\ \nabla_{\lambda} \mathcal{L}(\mathbf{x}, \lambda) = \mathbf{0} &\Rightarrow \mathbf{w}_\alpha^T \mathbf{x} = 1 \Rightarrow \lambda \|\mathbf{w}_\alpha\|_2^2 = 1 \\ &\Rightarrow \lambda = \frac{1}{\|\mathbf{w}_\alpha\|_2^2}, \\ \Rightarrow \mathbf{x} = \lambda \mathbf{w}_\alpha &= \frac{\mathbf{w}_\alpha}{\|\mathbf{w}_\alpha\|_2^2} = \frac{a^2 b^2}{a^2 \sin^2 \alpha + b^2 \cos^2 \alpha} \\ &\quad \times \begin{bmatrix} \cos \alpha & \sin \alpha \\ a & b \end{bmatrix}^T, \\ \Rightarrow \rho_\alpha = \|\mathbf{x}\|_2 &= \frac{1}{\|\mathbf{w}_\alpha\|_2} = \frac{ab}{\sqrt{a^2 \sin^2 \alpha + b^2 \cos^2 \alpha}}. \end{aligned} \tag{A3}$$

We now have ρ_α in Eq. (A3), but our goal is to eliminate α and obtain $\rho(\theta)$. From Eq. (A1), we know that lines that are orthogonal to the tangent line have a slope of $a^2 y/b^2 x$, so

$$\begin{aligned} \tan \theta &= \frac{a \sin \alpha}{b \cos \alpha} = \frac{a}{b} \tan \alpha, \\ \Rightarrow \cos^2 \alpha &= \frac{1}{1 + \tan^2 \alpha} = \frac{1}{1 + \frac{b^2}{a^2} \tan^2 \theta}, \end{aligned} \tag{A4}$$

$$\sin^2 \alpha = \frac{\frac{b^2}{a^2} \tan^2 \theta}{1 + \frac{b^2}{a^2} \tan^2 \theta}. \tag{A5}$$

Therefore, for a given θ , the corresponding $\rho(\theta)$ is obtained after substitution as

$$\rho(\theta) = \frac{ab}{\sqrt{a^2 \sin^2 \alpha + b^2 \cos^2 \alpha}} = \sqrt{a^2 \cos^2 \theta + b^2 \sin^2 \theta}. \tag{A6}$$

Equation (A6) allows us to obtain the set of points $\{(\rho, \theta)\}$ in the COTANS domain for a standard ellipse, but we wish to obtain such points for a general ellipse whose foci are \mathbf{p}_e and \mathbf{p}_r . We are given that the LOS distance between the emitter and receiver is d_{LOS} and that the distance of the NLOS reflection from a given boundary is d_{NLOS} . Our method is to first obtain (ρ, θ) for a standard ellipse with foci at $\mathbf{p}'_r = [-d_{\text{LOS}}/2 \ 0]^T$ and $\mathbf{p}'_e = [d_{\text{LOS}}/2 \ 0]^T$, with axes $a = d_{\text{NLOS}}/2$, $b = \sqrt{d_{\text{NLOS}}^2 - d_{\text{LOS}}^2}/2$. One example of such an ellipse is given in Fig. 8(a). The ellipse corresponding to the true localization scenario is obtained by rotating this standard ellipse as in Fig. 8(b), and then translating it to its correct position as in Fig. 8(c). The (ρ, θ) of the tangent line must be correspondingly transformed to match the new position of the rotated and translated tangent that yields the desired $(\rho_{\text{COTANS}}, \theta_{\text{COTANS}})$.

To modify (ρ, θ) as per Fig. 8, first consider the azimuth of the vector $\mathbf{p}_e - \mathbf{p}_r$ as θ_{rot} . To align the standard ellipse with the target ellipse, we replace each (ρ, θ) with $(\rho, (\theta + \theta_{\text{rot}}) \bmod 2\pi)$, which we term (ρ, θ'') . This rotation leaves the ρ -value of each tangent line unchanged, and only affects the azimuth, while the new foci are at \mathbf{p}''_e and \mathbf{p}''_r . Next, we calculate a translation vector $\mathbf{p}_{\text{trans}} = \mathbf{p}_r - \mathbf{p}''_r$, which would be added to any point on the rotated standard ellipse to obtain the target ellipse. To obtain the resulting $(\rho_{\text{COTANS}}, \theta_{\text{COTANS}})$ pairs, we first calculate the dot product $\rho_{\text{proj}} = \mathbf{p}_{\text{trans}} \cdot \hat{\rho}$, where $\hat{\rho} = [\cos \theta'' \ \sin \theta'']^T$ is the unit vector pointing towards the tangent line. Thus, we project the translation vector $\mathbf{p}_{\text{trans}}$ onto $\hat{\rho}$. If $\rho_{\text{proj}} \geq 0$, then we replace (ρ, θ'') with $(\rho + \rho_{\text{proj}}, \theta'')$. If $\rho_{\text{proj}} < 0$, then we replace ρ with $|\rho - |\rho_{\text{proj}}||$. If $|\rho_{\text{proj}}| < \rho$ and $\rho_{\text{proj}} < 0$, then we do not modify the azimuth θ'' ; else, we replace θ'' with $(\theta'' + \pi) \bmod 2\pi$. Carrying out these operations for each of the starting $\{(\rho, \theta)\}$ points, we obtain a final transformed set of points as $\{(\rho_{\text{COTANS}}, \theta_{\text{COTANS}})\}$.

2. COTANS transform for tangent planes to spheroids in 3D

The equation of an origin-centered spheroid is $(x^2/a^2) + (y^2/b^2) + (z^2/c^2) = 1$. A point (x_0, y_0, z_0) on its

surface can be parametrized as $(a \cos \alpha \sin \beta, b \sin \alpha \sin \beta, c \cos \beta)$, where α and β are intermediate parameters. Similarly to 2D, the tangent plane at (x_0, y_0, z_0) is $(xx_0/a^2) + (yy_0/b^2) + (zz_0/c^2) = 1$, and we have an optimization similar to Eq. (A2), this time using $\mathbf{w}_{\alpha,\beta}^T = [x_0/a^2 \ y_0/b^2 \ z_0/c^2] = [\cos \alpha \sin \beta/a \ \sin \alpha \sin \beta/b \ \cos \beta/c]$ instead of \mathbf{w}_α^T . A Lagrangian method as in Eq. (A3) yields

$$\rho_{\alpha,\beta} = \frac{1}{\|\mathbf{w}_{\alpha,\beta}\|_2} = \frac{1}{\sqrt{\frac{\cos^2 \alpha \sin^2 \beta}{a^2} + \frac{\sin^2 \alpha \sin^2 \beta}{b^2} + \frac{\cos^2 \beta}{c^2}}}, \tag{A7}$$

$$\mathbf{x} = \frac{\mathbf{w}_{\alpha,\beta}}{\|\mathbf{w}_{\alpha,\beta}\|_2^2} = \rho_{\alpha,\beta}^2 \begin{bmatrix} \cos \alpha \sin \beta & \sin \alpha \sin \beta & \cos \beta \end{bmatrix}^T. \tag{A8}$$

To eliminate α and β and express \mathbf{x} in terms of the azimuth θ and elevation ϕ , we use

$$\mathbf{x} = \rho_{\alpha,\beta} [\cos \theta \sin \phi \ \sin \theta \sin \phi \ \cos \phi]^T. \tag{A9}$$

First, noting that $\tan \phi = (a/b) \tan \alpha$, we have the same results of Eq. (A4) for $\cos^2 \alpha$ and Eq. (A5) for $\sin^2 \alpha$. Note that for the case of a spheroid where $a=b$, we have $\alpha = \theta$, since the x - y cross section of the ellipsoid is then a circle. To eliminate β , we observe that

$$\begin{aligned} \tan^2 \phi &= \frac{c^2 \cos^2 \alpha}{a^2 \cos^2 \theta} \tan^2 \beta = \frac{c^2}{a^2} \tan^2 \beta, \\ \Rightarrow \cos^2 \beta &= \frac{1}{1 + \frac{a^2}{c^2} \tan^2 \phi}, \end{aligned} \tag{A10}$$

$$\sin^2 \beta = \frac{\frac{a^2}{c^2} \tan^2 \phi}{1 + \frac{a^2}{c^2} \tan^2 \phi}. \tag{A11}$$

We substitute into Eq. (A7) our results in Eq. (A4), (A5), (A10), and (A11) to get

$$\rho(\theta, \phi) = \sqrt{c^2 \cos^2 \phi + a^2 \sin^2 \phi}. \tag{A12}$$

Note that because of symmetry around the z -axis, there is no θ -dependence at this stage.

Using Eq. (A12), we can take the COTANS transform of tangent planes to a standard spheroid with $a=b$. We can rotate this spheroid and translate it, and correspondingly transform the points (ρ, θ, ϕ) that define the tangent planes. Our operations are similar to those in the 2D case. We first define our standard spheroid to have foci at $\mathbf{p}'_f = [0 \ 0 \ -d_{\text{LOS}}/2]^T$ and $\mathbf{p}'_e = [0 \ 0 \ d_{\text{LOS}}/2]^T$, with axes $a=b = \sqrt{d_{\text{NLOS}}^2 - d_{\text{LOS}}^2}/2$ and $c = d_{\text{NLOS}}/2$. We calculate the azimuth and elevation of the vector $\mathbf{p}_e - \mathbf{p}_f$, and call these angles θ_{rot} and ϕ_{rot} , respectively. We rotate the spheroid by

ϕ_{rot} with the y -axis as the axis of rotation (thereby obtaining the correct elevation), then rotate the result by θ_{rot} with the z -axis as the axis of rotation (thereby obtaining the correct azimuth), and finally, we translate it to its correct position.

Rotating the spheroid by ϕ_{rot} will not change the ρ of a given (ρ, θ, ϕ) tangent plane, but may affect its azimuth and elevation. The unit vector orthogonal to (ρ, θ, ϕ) in Cartesian coordinates is given by $\hat{\mathbf{p}} = [\cos \theta \sin \phi \ \sin \theta \sin \phi \ \cos \phi]^T$. To rotate this vector with the y -axis as the axis of rotation, we multiply it by the corresponding rotation matrix to obtain

$$\begin{aligned} &\begin{bmatrix} \cos \phi_{\text{rot}} & 0 & \sin \phi_{\text{rot}} \\ 0 & 1 & 0 \\ -\sin \phi_{\text{rot}} & 0 & \cos \phi_{\text{rot}} \end{bmatrix} \begin{bmatrix} \cos \theta \sin \phi \\ \sin \theta \sin \phi \\ \cos \phi \end{bmatrix} \\ &= \begin{bmatrix} \cos \phi_{\text{rot}} \cos \theta \sin \phi + \sin \phi_{\text{rot}} \cos \phi \\ \sin \theta \sin \phi \\ -\sin \phi_{\text{rot}} \cos \theta \sin \phi + \cos \phi_{\text{rot}} \cos \phi \end{bmatrix}. \end{aligned} \tag{A13}$$

The azimuth and elevation of the vector obtained in Eq. (A13) are then substituted as the new θ and ϕ values of the plane, θ'' and ϕ'' . Next, we rotate the tangent plane by θ_{rot} , which does not change ρ or ϕ'' but replaces θ'' with $(\theta'' + \theta_{\text{rot}}) \bmod 2\pi = \theta'''$. At this stage, the point corresponding to the transformed value of \mathbf{p}'_f is a point \mathbf{p}''_f . Finally, we calculate the translation vector $\mathbf{p}_{\text{trans}} = \mathbf{p}_r - \mathbf{p}''_f$, and obtain $\rho_{\text{proj}} = \mathbf{p}_{\text{trans}} \cdot \hat{\mathbf{p}}$. The transformation of the tangent plane in the COTANS domain is similar to the 2D case. If $\rho_{\text{proj}} \geq 0$, we replace $(\rho, \theta''', \phi''')$ with $(\rho + \rho_{\text{proj}}, \theta''', \phi''')$. If $\rho_{\text{proj}} < 0$, then we subtract the projection result from ρ , thus replacing ρ with $|\rho - |\rho_{\text{proj}}||$. If $|\rho_{\text{proj}}| < \rho$ and $\rho_{\text{proj}} < 0$, then we do not modify the azimuth θ''' or the elevation ϕ''' ; else, we replace θ''' with $(\theta''' + \pi) \bmod 2\pi$ and ϕ''' with $\pi - \phi'''$. Carrying out this procedure for each of the starting $\{(\rho, \theta, \phi)\}$ points, we obtain a final set of $\{(\rho_{\text{COTANS}}, \theta_{\text{COTANS}}, \phi_{\text{COTANS}})\}$ points in the 3D COTANS domain.

APPENDIX B: PSO IMPLEMENTATION

In the standard PSO algorithm (Clerc, 2010), we define N particles $\mathbf{x}_i^{(k)}$, corresponding to points $\mathbf{x}_i^{(k)} = [x_{i,1}^{(k)} \dots x_{i,D}^{(k)}]^T$ in a D -dimensional search space at iteration k of the PSO algorithm. This search space is delimited by the boundaries $[x_{j,\text{min}}, x_{j,\text{max}}]$ for each of the $j \in \{1 \dots D\}$ parameters that we optimize over. Each particle has a velocity vector $\mathbf{v}_i^{(k)} = [v_{i,1}^{(k)} \dots v_{i,D}^{(k)}]^T$, assigned randomly at initialization. Particles remember their best locations \mathbf{p}_i in the search space over iterations 1 through k , where their objective function $f(\mathbf{p}_i)$ was the largest. Finally, each particle is in a set of M randomly chosen neighbors, over which we determine the best location $\mathbf{p}_{i,g}$ over past iterations. The PSO update step for the velocity and position of each particle is then given for each dimension d by

$$v_{id}^{(k+1)} \leftarrow \alpha v_{id}^{(k)} + U(0, \beta) (p_{id} - x_{id}^{(k)}) + U(0, \beta) (p_{ig} - x_{id}^{(k)}), \quad (\text{B1})$$

$$x_{id}^{(k+1)} \leftarrow x_{id}^{(k)} + v_{id}^{(k+1)}, \quad (\text{B2})$$

where α and β that define a particle's inertia and acceleration for updating its velocity (standard values are 0.7298 and 1.4986, respectively), and $U(0, \beta)$ is a sample of the uniform random variable. Particles that exit the boundaries of the search space are reflected back.

¹In practice, if $T_p > T_c$, then T_p can be estimated, but this is an environment-specific problem outside the scope of our method, as is an estimation of T_m .

²In our scale watertank setting, the regularly-spaced grid-squares have length 1 mm.

³More sophisticated methods for extracting the LOS peak could also be used (Guvenc and Sahinoglu, 2005).

⁴Algorithm 2 does not require leveraging all of the boundaries to yield an accurate result. The additional NLOS arrivals provide spatial diversity that helps to increase localization accuracy and mitigate occlusion.

⁵We use a resolution of 0.1 mm for ρ and 0.1° for θ and ϕ in our experiments.

⁶We heuristically use a uniform averaging filter that is 3 mm wide in ρ , and 3° wide in θ and ϕ .

⁷Given a maximum of the accumulator result at some $(\rho_{\max}, \theta_{\max}, \phi_{\max})$, we zero out the θ and ϕ within $\pm 15^\circ$ of the maximum, and all ρ values within this sector.

⁸70% below the predicted magnitude, in our implementation.

Alameda-Pineda, X., and Horaud, R. (2014). "A geometric approach to sound source localization from time-delay estimates," *IEEE/ACM Trans. Audio. Speech. Lang. Process.* **22**(6), 1082–1095.

Baggeroer, A. B., Kuperman, W. A., and Mikhalevsky, P. N. (1993). "An overview of matched field methods in ocean acoustics," *IEEE J. Oceanic Eng.* **18**(4), 401–424.

Borrmann, D., Elseberg, J., Lingemann, K., and Nüchter, A. (2011). "The 3D Hough transform for plane detection in point clouds: A review and a new accumulator design," *3D Res.* **2**(2), 3.

Brutti, A., Omologo, M., and Svaizer, P. (2010). "Multiple source localization based on acoustic map de-emphasis," *EURASIP J. Audio Speech Music Process.* **2010**, 1–17.

Chitre, M. (2007). "A high-frequency warm shallow water acoustic communications channel model and measurements," *J. Acoust. Soc. Am.* **122**(5), 2580–2586.

Chitre, M., and Pelekanakis, K. (2014). "Channel variability measurements in an underwater acoustic network," in *Proceedings of UComms 2014*, September 3–5, Levante, Italy, pp. 1–4.

Clerc, M. (2010). *Particle Swarm Optimization* (John Wiley & Sons, New York).

Deane, G. B. (1994). "A three-dimensional analysis of sound propagation in faceted geometries," *J. Acoust. Soc. Am.* **96**(5), 2897–2907.

Demirli, R., and Saniie, J. (2001). "Model-based estimation of ultrasonic echoes. Part I: Analysis and algorithms," *IEEE Trans. Ultrason. Ferroelect. Freq. Contr.* **48**(3), 787–802.

Dokmanic, I., Parhizkar, R., Ranieri, J., and Vetterli, M. (2015). "Euclidean distance matrices: Essential theory, algorithms, and applications," *IEEE Signal Process. Mag.* **32**(6), 12–30.

Fischell, E. M., Kroo, A. R., and O'Neill, B. W. (2019). "Single-hydrophone low-cost underwater vehicle swarming," *IEEE Robot. Autom. Lett.* **5**(2), 354–361.

Fleury, B. H., Tschudin, M., Heddergott, R., Dahlhaus, D., and I. Pedersen, K. (1999). "Channel parameter estimation in mobile radio environments

using the SAGE algorithm," *IEEE J. Select. Areas Commun.* **17**(3), 434–450.

Guvenc, I., and Sahinoglu, Z. (2005). "Threshold selection for UWB TOA estimation based on kurtosis analysis," *IEEE Commun. Lett.* **9**(12), 1025–1027.

Hermard, J.-P. (1999). "Broad-band geoacoustic inversion in shallow water from waveguide impulse response measurements on a single hydrophone: Theory and experimental results," *IEEE J. Oceanic Eng.* **24**(1), 41–66.

Kim, J., Kim, J., Nguyen, L. T., Shim, B., and Hong, W. (2019). "Tonal signal detection in passive sonar systems using atomic norm minimization," *EURASIP J. Adv. Signal Process.* **2019**, 43.

Knapp, C., and Carter, G. (1976). "The generalized correlation method for estimation of time delay," *IEEE Trans. Acoust. Speech, Signal Process.* **24**(4), 320–327.

Korhonen, T. (2008). "Acoustic localization using reverberation with virtual microphones," in *Proceedings of the IWAENC 2008*, September 14–17, Seattle, WA, pp. 211–223.

Le Gall, Y., Socheleau, F., and Bonnel, J. (2016). "Matched-field performance prediction with model mismatch," *IEEE Signal Process. Lett.* **23**(4), 409–413.

Li, L., and Krolik, J. (2013). "Simultaneous target and multipath positioning," *IEEE J. Sel. Top. Signal Process.* **8**(1), 153–165.

Naseri, H., Costa, M., and Koivunen, V. (2014). "Multipath-aided cooperative network localization using convex optimization," in *Proceedings of the 48th Asilomar Conference on Signals, Systems and Computers*, November 2–5, Pacific Grove, CA, pp. 1515–1520.

Naseri, H., and Koivunen, V. (2016). "Cooperative simultaneous localization and mapping by exploiting multipath propagation," *IEEE Signal Process. Mag.* **65**(1), 200–211.

Qiao, T., Redfield, S., Abbasi, A., Su, Z., and Liu, H. (2013). "Robust coarse position estimation for TDOA localization," *IEEE Wireless Commun. Lett.* **2**(6), 623–626.

Ribeiro, F., Zhang, C., Florêncio, D. A., and Ba, D. E. (2010). "Using reverberation to improve range and elevation discrimination for small array sound source localization," *IEEE Trans. Audio. Speech. Lang. Process.* **18**(7), 1781–1792.

Salmi, J., Richter, A., and Koivunen, V. (2008). "Detection and tracking of MIMO propagation path parameters using state-space approach," *IEEE Trans. Signal Process.* **57**(4), 1538–1550.

Scheuing, J., and Yang, B. (2008). "Disambiguation of TDOA estimation for multiple sources in reverberant environments," *IEEE Trans. Audio. Speech. Lang. Process.* **16**(8), 1479–1489.

Setlur, P., Smith, G. E., Ahmad, F., and Amin, M. G. (2012). "Target localization with a single sensor via multipath exploitation," *IEEE Trans. Aerosp. Electron. Syst.* **48**(3), 1996–2014.

Srinivasan, P., Liang, P., and Hackwood, S. (1990). "Computational geometric methods in volumetric intersection for 3D reconstruction," *Pattern Recogn.* **23**(8), 843–857.

Su, Z., Shao, G., and Liu, H. (2017). "Semidefinite programming for NLOS error mitigation in TDOA localization," *IEEE Commun. Lett.* **22**(7), 1430–1433.

Sun, W., and Wang, Z. (2018). "Online modeling and prediction of the large-scale temporal variation in underwater acoustic communication channels," *IEEE Access* **6**, 73984–74002.

Too, Y. M., Chitre, M., Barbastathis, G., and Pallayil, V. (2017). "Localizing snapping shrimp noise using a small-aperture array," *IEEE J. Oceanic Eng.* **44**(1), 207–219.

Venkateswaran, S., and Madhow, U. (2012). "Localizing multiple events using times of arrival: A parallelized, hierarchical approach to the association problem," *IEEE Trans. Signal Process.* **60**(10), 5464–5477.

Waterston, J., Rhea, J., Peterson, S., Bolick, L., Ayers, J., and Ellen, J. (2019). "Ocean of Things: Affordable maritime sensors with scalable analysis," in *OCEANS 2019-Marseille*, pp. 1–6.

Xu, W., Baggeroer, A. B., and Schmidt, H. (2006). "Performance analysis for matched-field source localization: Simulations and experimental results," *IEEE J. Oceanic Eng.* **31**(2), 325–344.