

Multi-Robot Teamwork in Multistatic Sonar

F. Ehlers, *Member, IEEE*

Abstract—In this paper we present a novel and relevant application for multi-robot systems which is to implement autonomy for Multistatic Sonar. Multistatic Sonar is used in Antisubmarine Warfare and exploits multiple aspects on the target generated by multiple acoustic sources and receivers. We propose to use a team of unmanned underwater vehicles starting with search operations as covert receivers and then switching to pursuit operations as actively communicating multistatic receivers. With the help of a simulation setup and data from recent experiments at sea, we explain what are the specific challenges related to the Multistatic Sonar setup. These challenges are mainly related to the measurement model and to communication constraints. We then briefly outline a potential “baseline” solution, presented as a survey, for the overall system architecture capable to face these challenges.

I. INTRODUCTION

THE general problem tackled in the paper is that of decentralized sonar detection and tracking with a team of autonomous agents (typically AUVs, Autonomous Underwater Vehicles) [1]. The problem is further made more complex by communication constraints: Underwater acoustic communication is notoriously affected by bandwidth and range limitations; in addition, the team members wish to remain covert, and hence they do not want to communicate to each other, unless to hand over to the other team members not only the information but also the tracking task. They have to do so for two reasons: (i) because only the diversity generated by multiple aspects on the target allows a sufficiently high probability of detection [2] and (ii) because the target can be much faster than a single AUV.

This so-called ‘Collaborative Antisubmarine Warfare’ [3] follows the natural trend in surveillance applications to incorporate distributed sensors. This is in particular true for the multistatic sonar outlined above and for multistatic radar. In both applications different aspect angles on hypothetical targets are exploited by spatially separated sensors. And it is true for multi-sensor applications aiming at a higher level of situational awareness where physically independent (and desirably quasi complementary) sensor types are used. Hence, findings from working on the Multistatic Sonar problem are also important for a large variety of other surveillance applications.

Adding ‘autonomy’ to the multi-sensor design is important because it offers added-value in terms of persistence, cost reduction and risk minimization. Adding ‘autonomy’ to

multistatic sonar has to face especially the challenge that ‘teamwork’ is required between team members including the coordination of work schedules on potential targets (including a lot of clutter targets) while communication between team members can only happen on a very limited basis. In other words, the autonomous multistatic sonar will fail if the AUVs show pathological behavior, e.g. they all stick to the same target (which might not even be the real target, but a false alarm) and it will fail if they all communicate (indicating where they are and allowing the target to make a plan to evade). We propose to overcome this challenge by introducing a learning phase especially for teamwork. In team sports, team members are exploring the value of collaboration during training sessions. For example in basketball, they exercise on the so-called ‘no-look pass’ between them. We make the analogy between ‘no-look’ and no-communication and propose a training (or learning) phase for the specific autonomous teamwork features. We explain how this analogy influences the overall design architecture of the multistatic surveillance team.

The paper has two main contributions:

- Advertising the Multistatic Sonar task as a suitable playing field for multi-agent surveillance systems and as an example to cross-check how much it fits into a potential unified formulation.
- Implementation of a hierarchical structure to separate the learning processes related to teamwork from the data driven control and planning of the individual agents.

The paper is organized as follows: In the next section, we describe the Multistatic Sonar system. In section III, a simulation environment is proposed. With the help of this simulation task, we point to open research tasks. Besides the planning and decision making process on board each individual AUV, an interesting new research task is on the implementation of team behavior. In section IV we briefly outline a “baseline” solution by listing tools available in the robotics or autonomy literature. In particular, we propose a reinforcement learning algorithm for the teamwork. We then show how this leads to methods related to the Potts model, which is well known in statistical physics and can be used to derive a solution in the form of an “activity map” for the AUVs. Conclusions are given in section V.

Manuscript received March 5, 2010.

F. Ehlers is with the NATO Undersea Research Centre (NURC), La Spezia, Italy (phone: 0039-0187-527416; fax: 0039-0187-527330; e-mail: ehlers@nurc.nato.int).

II. DESCRIPTION OF A MULTISTATIC SONAR SYSTEM

In this section, we introduce the multistatic measurement system. The description summarizes state-of-the-art technology for manned systems and hence serves as the basis for plans on simulations and experiments executed in order to add autonomy to this surveillance system.

A. Teamwork in Multistatic Sonar

Multistatic active sonar involves multiple entities transmitting signals and receiving echoes. Receivers can be kept covert if they are spatially separated from the transmitter. Multistatic sonar is a coordinated effort aimed at gaining a more accurate view of the target and discriminating it from the background. However, multistatic active sonar is more than sonar technology. It is a system of hardware platforms and sophisticated software modeling and signal processing algorithms that are networked together. Especially when only underwater communications is available and, for tactical reasons, covert receivers have to be used, the resulting limited network capabilities are a constraint in the sense that decisions have to be made on which data should be exchanged. The envisioned autonomous multistatic system is scalable such that it can be expanded depending on its mission and scope.

B. Advantages of Multistatic Sonar

Multistatic active sonar has the obvious benefit of combining data from multiple pairs of sonar transmitters and receivers. In other words, there are more “ears” in the water to improve the detection, localization, and identification of undersea objects. This results in a reduced false alarm rate and the estimation of the target state vector (location, heading & speed) is of higher quality and accuracy, provided the (at least partially) fusion of gathered “single ear” information.

When using multiple sources, a single receiver can already exploit these “data fusion benefits” of different bistatic aspects, without communicating, i.e. it can operate completely covert. This covertness does not allow the target acting cleverly, thus its stealth capability diminishes. The cost of losing covertness can be quantified by comparing the

detection performance for two different Target Strength values: assigning to the clever target the lowest value and for the innocent target a substantially higher value.

C. Physical agents carrying sonar receivers

Fixed and moving assets can be mixed. At NURC, a transmitter buoy is available (Fig. 1, left) and as receivers Autonomous Underwater Vehicles (Fig. 1, right) are used towing a receiving array and a quite powerful transmitter (given limited battery capacity) that is proposed in our context to be used for reliable long-range communication purposes. The source buoys can also be used for underwater communication (given an adequate receiving unit), and they can receive/exchange commands & data via radio (allowing the utilization of communication gateway buoys). The AUVs have low speed compared to the speed of the target.

D. Multistatic data fusion architecture

The software behind multistatic acoustic systems includes signal processing algorithms referred to in the field as data fusion. Data fusion merges the data taken from the different tools in the multistatic system into a coherent picture. This picture reveals a series of contact points, which create a track showing the target’s progress through the water.

1. The signal processing onboard each AUV includes beamforming, matched filtering and detection. It is generating a list of contacts for each transmission (also called a “ping”) from an active source.
2. Fusion and tracking for all data available onboard the AUV is performed. There are no bandwidth limitations onboard each AUV. Hence, a centralized tracking version will allow associations (and probabilities for being correct associations) between the list of contacts from all different source receiver pairs. In this step, multistatic advantages in detection and localization performances are locally exploited as much as possible.
3. Tracking results are distributed between the receivers. The networking capabilities introduce constraints on how many tracks can be exchanged between the AUVs.
4. With the additional information from other receivers, data fusion and updated tracking, algorithms are able to further exploit the advantages of the multistatic setup, i.e. resulting in smaller uncertainty covariance matrices, allowing smaller association gates in the tracking, and by that, finally, leading to a lower false track rate.

As written in step 2) and 4) tracking is a major part of the multistatic software: In fact, previous research has shown that tracking algorithms like multi-hypotheses tracking (MHT) not only provide a very elegant framework to perform this task, but that even more the “integration of information” over time is necessary to actually do the association, i.e. kinematic information is necessary to associate contacts in the stochastically fluctuating sonar pictures [4].

Sonar performance prediction modeling shows that only rarely do distributed sensors have a similar quality on a

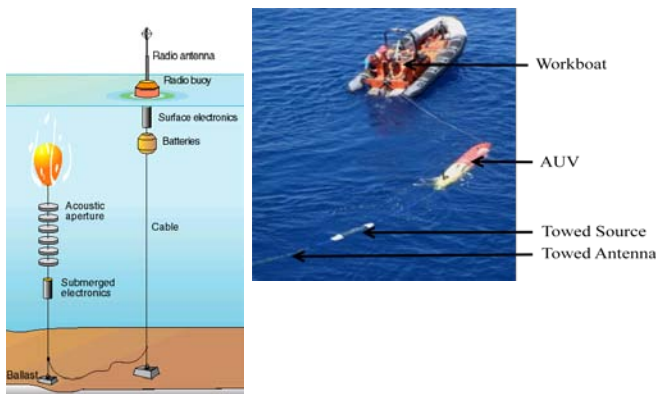


Figure 1: (left) Stationary acoustic source, (right) Deployment of the OEX Explorer AUV and its towed source, together with a slim towed array. (The workboat is only necessary for deployment and recovery.)

specific ping of the target track [5]. Therefore, data fusion algorithms must be based on realistic modeling of sensor performance for each sensor, each ping, and each target. Data is weighted so that data which is considered more accurate or valid is given more weight in the algorithms. Since in an autonomous application, receivers are able to plan their actions, the same modeling of data quality that is used to calculate the weights for data fusion can be used in prediction and path planning. Furthermore, as simulations will play a major role as part of the learning process for the team specific decision making process, these models will also be used when running the ‘training simulation’ sessions. In the proposed autonomous framework, the steps 3) and 4) directly lead to questions of decision making and teamwork:

For the AUV, the cost of communication is to lose covertness which has to be quantified as losing detection opportunities. On the other hand, the best state estimation is possible, when receivers communicate: This is not only true for the multistatic fusion of gathered data, but even more inevitable when looking at the path planning for the slow AUVs that have to follow a fast target. In this case the AUVs have to organize a kind of hand-over of target information. The output consists of a set of tracks that is updated in real-time as a new contact is received. A contact contains bearing and receiving time information. The tracking is implementing Bayesian updates to estimate the target state. Hence, each track includes the following information: the position of the target, the speed and course of the target, track localization accuracy, and track likelihood. Instead of the display of the output in manned systems, in the autonomous world, the output has to lead to actions automatically. The linkage between measurement and actions is further discussed in section III.B.

E. Real data

In 2009, NURC together with its partners has conducted sea trials using active sources and a receiver as described before. Of course, the sea trial in this early stage of technology could not focus on the implementation of different autonomous team controls, but already relevant data sets have been gathered that demonstrate the relevance of the research proposed in the previous sections.

As an example, Fig. 2 shows the geo-referenced output of the signal processing onboard the AUV system. Two stand-off sources were ensonifying a cluttered area. Each pixel in Fig. 2 is representing the energy reflected from a 25m x 25m cell. The figure does not show a target echo, just reverberation clutter. A target echo would (as a single contact without kinematic information) not be different from these echoes. The picture shows “cross-fixed” bistatic contacts (demonstrating the triangulation capability of multistatics) belonging to two different source-receiver-geometries. But also, (geometry dependent) echoes are shown for which a cross-fixing is not possible. Performing the correct association is the challenge in these shallow water areas, and tracking has been proven to counter this challenge.




Figure 2. Geo-referenced output of the signal processing of echoes from two bistatic sources recorded on NURC’s AUV during GLINT 09. Each pixel is representing the energy reflected from a 25m x 25m cell. Clutter returns are visible as colored lines (parts of an ellipse defined by the positions of the acoustic source and the receiver). Two differently oriented lines are shown, referring to two different (spatially separated) sources. The receiver simultaneously listens to echoes from the two different sources. Multistatic data fusion (cross-fixing) can be applied if echoes are detected for both geometries.

Another valuable input from the experiment for the simulation in the test bed is that AUVs (although towing an array) can manage a very high turn rate. This maneuverability makes the proposed target-adaptive control become realistic.

Underwater communication with the robust JANUS standard has achieved a data rate sufficiently high for the described purpose with a very high reliability at relevant distances in the recent experiment SUBNET 09 [6].

In summary, hardware and software exists to implement multistatics up to data fusion and tracking. What is missing, and what is a logical follow-on, is to add autonomy.

III. PROPOSED SIMULATION ENVIRONMENT TO TEST TEAM ARCHITECTURES

The objective of the simulation environment is the comparison of multiple approaches and techniques in using multiple AUVs equipped with receiver arrays to detect and track a target. In this paper, we are leaving out questions like for example path planning of the AUVs, robustness of tracking, number of AUVs needed, the effect of the target speed on the tracking performance. Discussions on these and related relevant and active research fields can be found in [1,4,5]. Instead, we focus on the overall architecture for the ‘system of systems’. In the next subsection, we describe the simulation scenario. Then, we outline the overall system structure and focus on the role of teamwork and the issues related to it.

A. Simulation scenario

We propose to simulate a multistatic survey and pursuit in clutter. The geometry of the test bed is depicted in Fig. 3. (Note: The region depicted has just been chosen as an example. As this is a shallow water region with known environmental acoustic parameters (e.g. sound speed as a function of geographical location, depth, season) and known

bathymetry, it would be well suited to establish benchmark tests. Furthermore, a final demonstration at sea could be organized/planned/executed in this area.) Two stand-off sources ensonify (“illuminate”) a shallow water area. Since sound is reflected from a complex sea bottom structure, false alarms are generated by these clutter returns. The best way to overcome this is to send AUVs, equipped with towed array into the area, using their good angular resolution and, hence, their ability to discriminate target echoes from background echoes when they are close enough. As a target, a small submarine is hiding in this area. Its evasive action is to move to North East. The task of N (in Fig. 3: $N=16$) AUVs is to organize themselves in a way that they keep the most accurate track of the target under the constraints of both: a minimal usage of communication and a minimal number of AUVs involved in the tracking task. Minimal communication is important, because covertness is a big advantage. Involving a minimal number of AUVs is important to prevent being exhausted by false alarms/false tracks (which cannot be avoided in difficult measurement environments). Practically, communication from the AUVs to the stand-off sources could be envisioned, enabling the idea to implement ways of directional/potentially covert communications. The stand-off sources take the role of a communication gateway, distributing information. They could also be seen to be a centralized planning unit (giving advice to the AUVs comparable to a coach of a soccer team during a game). Because they do not know the positions of the AUVs that remain covert they have to do the planning respecting this uncertainty and the final decision has to be made inside each covert AUV. It will be important to model the target as a clever one: As soon as it knows the positions of an AUV (because the AUV is communicating), it avoids it. In this simulation, we are able to model state-of-the-art signal processing and data fusion [5].

B. The overall system architecture

We can formally write the solution of the control problem for the envisioned simulation scenario following the description as a stochastic nonlinear minimax dynamic game with noisy measurements [8]. However, we will not be able to solve it with standard methods (even numerically) because of its complexity. For a particle based numerical solution, a decreased complexity would lead to a feasible number of particles needed. Thinking about this particle stream like simulation, it is shown in [9] that taking out particles associated with low probability greatly minimizes the demand on processing power in the numerical simulation. We interpret “taking out particles” in our context as “making decisions”, and by introducing a reasonable structure for the decision making process, we generate a computational feasible solution of the formally written control problem. (Note: In this paper, we give a simplified description of how the individual linkage between measurement and action looks like, because we want to focus for the remainder of the paper on the teamwork aspect of the multistatic surveillance tasks. However, we see this

simplified description to be conceptual rich enough to be expanded to incorporate all details of a real-world application.)

The decision making process envisioned for the multistatic surveillance task should be data driven, because we know from previous experiments (as described in section II), that the quality of the achievable data fusion depends on the quality of the contributing sensors which then depend on the hypothetical target track itself. This falls into the framework of partially observable robot control problems, in particular of the partially observable Markov decision process (POMDP). In order to keep the overall complexity of the decision making process low, we separate the individual path planning (data driven because it depends on the specific target track) from the team behavior planning where the dependence on the actual data is interpreted in a much wider sense: For each of a few “major direction for evasion” teamwork plans are generated. In Fig. 3 the major direction for evasion is Northeast, and we think that all together 8 directions (dividing 360 degrees into 45 degrees sections) are sufficient to describe the operational characteristics dependent on the submarine evasion process.

The individual decision level is modeled by the three output actions of a POMDP: (i) Pursue, (ii) Wait for more data before making a decision, (iii) Do not pursue. Since the ‘state of the world’ cannot be sensed directly and since there is the uncertainty in control effects, a POMDP an appropriate tool.

In case the individual POMDP action is ‘Pursue’, this action will not be executed if the trained teamwork plan does not allow the contribution of this specific AUV for the specific track at a specific position with a specific ‘major direction for evasion’.

(Note: The ‘major direction for evasion’ can either be predefined as in Fig. 3, or it can be estimated from the target track by linear regression of estimated target positions.)

The remaining question is now how to generate the ‘activation map’. A potential answer is given in the next section.

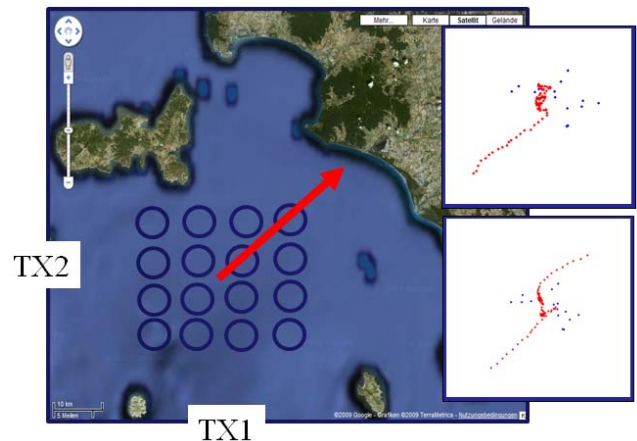


Fig.3 Test bed: Geometry in an exemplary environment (left). Simulation results of teamed AUVs trapping a target (right): For a slow target (top right) and a faster target (bottom right).

Table I summarizes the parameters that define the actual simulation.

Quantity	Interesting interval
Number of AUVs	2-32 (simulations Fig.3 N=16)
Initial patrol radius	1 km -5 km
Max speed of target	10 knts -30 knts
Max speed of AUVs	1 knts -4 knts
Surveillance radius	2 km -7 km
Close range (causing evasive action by the target)	0.5 km -1.5 km

IV. ARCHITECTURE IMPLEMENTING ADAPTIVE TEAM BEHAVIOR

In this section, we list available tools from the robotics or autonomy literature. We see these tools as building blocks that lead to a baseline solution, provided that the interfaces between them are appropriately designed.

A. Probabilistic multi-agent control

In order to implement multi-robot coordination the artificial intelligence community has a considerable interest in decentralized control of multiple decision makers or “agents” under uncertainty [10]. Formal models and algorithms have recently been developed. Decentralized schemes are a special demand in the underwater scenario because communication capabilities are physically limited. Because the probabilistic robotics approach aims at low computational load during execution and at robustness, it looks suitable to lead to a more practical way ahead.

The linkage between measurement and system analysis to the control architecture is implemented by choosing “critical system parameters” as input into the control procedures:

- 1) *Whether or not a target is observable (or reachable) for a team member and how far it is away from the at-best reachable curve. This concludes in a pessimistic view for each team member on whether or not its contribution to the hunting-team is possible.*
- 2) *In case this pessimistic view allows a participation in the hunting-team, path planning is needed. The path planning algorithm has to incorporate uncertainties on the target state, using the output of the tracking, see also paragraph IV.D.*
- 3) *However, in a team also other team members might achieve the same goal (maybe, even faster). Therefore, also a path planning with uncertainty for all other team members is performed. But, the complexity of an algorithm that incorporates the uncertainty for all other team members is too large. Hence, a-priori knowledge has to be generated. We propose to introduce a learning step for teamwork before the actual operational execution of the team surveillance. In this learning step, maps will be generated that show the activity of all AUVs*

given a specific ‘target scenario’. The target scenario is a coarse grained representation of the actual measurement. By this, a minimum of data-driven input is kept in the teamwork execution while allowing to exploit a-priori generated team knowledge, see also paragraph IV.E.

This describes the “sequencing part” of the test bed architecture where decisions are made to switch modes: From being covert in survey-mode to become a member of the hunting-team with active communication. This decision will be based on a policy that results from a cost minimization (or learning) process. The policy is literally a table linking beliefs (updated by measurements) to actions, but additionally covering robustness by following the probabilistic approach.

The advantages of this proposed approach are:

- The probabilistic approach is reflecting the overall stochastic nature of the problem.
- The approach is explicitly allowing multistatic tracking, which is a method proven to be able to provide sufficient performance even in areas with high false alarm rates.
- The approach is scalable in the sense that the number of assets is not explicitly set in its architecture.

The approach allows to model evasive behavior of the target in the “minimax” [8] control philosophy by implementing a cost maximization to calculate its movements.

B. Decentralized POMDP (DEC-POMDP)

Following [10], the DEC-POMDP model explicitly contains communication actions and assigns costs for each transmitted message. There is a definition for local policies for actions for an agent, which are a mapping from local histories of observations and histories of messages received to control actions. And also, there is a definition of a local policy for communication for an agent. Accordingly, a joint policy for a DEC-POMDP-COM is a tuple of local policies, one for each agent, where each of these is composed of the communication and action policies for each agent.

C. Path planning under uncertainty

Steps 2) and 3) in the proposed architecture IV.B) include the solution of path planning tasks. The formulation of the path planning has to be in extended space in order to include the uncertainties given in the target position. (For Step 3) also uncertainties on the position of the receiver have to be incorporated.) The initial uncertainties are calculated in the target tracking step. How the uncertainties will evolve in the future is answered using the prediction methodology from the Bayesian inference. It is important to maintain a high-fidelity sensor/measurement model. NURC has developed an advanced multistatic simulation tool, called MSTPA [5].

D. Discrete Input to POMDP

In order to implement the POMDP framework, measurements have to be discrete. Furthermore, to keep

models and associated learning procedures fast, we are proposing a very coarse grained representation for the input to the POMDP.

E. Teamwork

In this section, we specifically address the problem of preventing the pathological behavior where all AUVs pursue the same target track. We assume that a sufficient number of receivers for tracking a target is only the minority of all available receivers. Hence, a solution of this very simplified problem can be found if we recall the ‘minority problem’, well studied in game theory, as a motivating example. The main features of the ‘minority problem’ are: (i) Decisions by the team members are made independently, (ii) a reward is given after the decision has been made. These features fit exactly to the described multistatic teamwork problem: (i) The AUVs cannot communicate before making a decision, (ii) whether the overall task of tracking the target is performed efficiently, or not, will only be found after the decision has been made. In the next subsection we explain, how the ‘minority problem’ is solved by learning. In subsection IV.F, we refer to the analogy between the proposed learning scheme and the Potts model known in statistical physics. By simulating the Potts model, we are able to quickly generate ‘team activation maps’ that represent a learned solution for the ‘minority problem’ in an extrapolated form as it occurs in our application. The Potts model provides a linkage to statistically meaningful performance measures for the overall surveillance team, which then can be used to be optimized in the design process. This optimization process is then the ‘engineering paradigm’ for the described multi-agent system (or ‘system of systems’).

F. Learning as a solution for the ‘minority problem’

In order to overcome the problem of too many AUVs following the target, a similar but simplified problem is studied: The El Farol bar problem is a “minority game” [7]. It is described as follows: If too many people attend the bar, the patrons will not enjoy the evening. The problem for each potential visitor is to find out in advance whether the bar will be crowded below 50% or not. Therefore each of the N possible visitors I must decide whether to go or stay home. Only those agents who are in the minority, i.e. who went, when the bar was less than half full, or stood home when the bar was overcrowded, win. Decisions are made independently, i.e. the agents cannot communicate with each other directly. The only information available to them is the decision of the minority in the last L time-steps which determines the reward which they receive, after having made their decision.

In order to make the agents as clever as possible, each agent is modeled by an adaptive probabilistic Boolean function with L inputs. We assume that there is an odd number of N players, each of which receives a binary input. Each player I reacts to this input with a binary output. The way this reaction is generated is trained in a reinforcement manner by

strengthening weights for correctly behaving agents. It can be shown [7] that this approach maps the solution of the minority problem to different energy minima of an all-to-all coupled antiferromagnet.




Fig. 3. Output of a 16-states Potts model running on a 256 by 256 grid with uniform neighbor connections on a nearest neighbor version of the Metropolis algorithm. [source code available at <http://www.ensea.fr/applets/spin-models/montecarlo.zip>] The colors correspond to the activation area for a specific sensor. By tuning the external force, the neighborhood and initial condition, application adaptive maps can be generated quickly.

G. Physics-based learning by Potts Model

Extending the method in IV.F from a Boolean output to model “team activation maps” for multi-agents leads from the antiferromagnet model to the Potts model. The Potts model is normally simulated by the Metropolis algorithm, which allows to set initial conditions, neighborhood coupling and external forces. By controlling these parameters maps like in Fig. 3 can be generated, that represent the solutions of a repeated learning of weights procedure as referred to in the previous subsection. We propose to use available descriptive values for the Potts model (e.g. correlation length, critical exponents) also for the performance description of the multi-agent system.

V. CONCLUSION

Multistatic operations with autonomous underwater vehicles enable new anti-submarine warfare missions. An operationally relevant example is described as a test bed problem which allows research on various techniques for teamed distributed agents, all under the umbrella of game and control theory with a probabilistic multi-sensor setup for data fusion and tracking. As an appetizer for a lot of challenging tasks related to this simulation scenario, the problem of ‘learning the teamwork’ is addressed by a connection to physics-based learning mechanisms for multi-agent systems.

ACKNOWLEDGMENT

The author thanks the members of the GLINT 09 and SUBNET 09 teams, and K. McCoy, T. Curtin, J. Pitton, S. Kemna and R. Streit for initial discussions and support.

REFERENCES

- [1] F. Ehlers, “Autonomy in Multistatic Sonar,” Proceedings of the Machine Intelligence Workshop, Lerici, Italy, 2009.
- [2] F. Ehlers, “Particle filter for information management in an autonomously operating activated distributed sensor network,” Proceedings of OCEANS’07, Aberdeen, June 2007.

- [3] D. T. Hughes, R. Been, and A. Vermeij, "Heterogenous underwater networks for ASW: Technology and techniques," in *Maritime Systems and Technology*, Genova, Italy, 2008.
- [4] Ehlers, F., Daun, M., Ulmke, M., "System design and fusion techniques for multistatic active sonar," in *Proceedings of IEEE OCEANS '09*, art. no. 5278319, Bremen, 2009.
- [5] A. Wathelet, C. Strode, A. Vermeij, "Simulating Multistatic ASW Operations with the Multistatic Tactical Planning Aid (MSTPA) for Tactics Development," NATO RTO Modelling and Simulation Group Conference on 'Improving M&S Interoperability, Reuse and Efficiency in Support of Current and Future Forces', Prague, Czech Republic, on 4-5 October 2007.
- [6] K. McCoy,"JANUS: From Primitive Signal to Orthodox Networks," in *Proceedings of the Underwater Acoustics Measurements (UAM) 2009 Conference*, Nafplion, Greece, June 2009.
- [7] H.G. Schuster, *Complex Adaptive Systems*, Scator Verlag, 2001.
- [8] C. D. Charalambous, "Stochastic Nonlinear Minimax Dynamic Games With Noisy Measurements," *IEEE Trans. Automatic Control*, vol. 48, no. 2, Feb. 2003.
- [9] H. J. Kappen,"Path integrals and symmetry breaking for optimal control theory," *J. Stat. Mech.* (2005) P11011.
- [10] S. Seuken, S. Zilberstein, "Formal models and algorithms for decentralized decision making under uncertainty," *Auton. Agent Multi-Agent Syst.*, DOI 10.1007/s10458-007-9026-5.