LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

LLNL-TR-795739

17-ERD-101 LDRD Final Report

R. A. Goldhahn

October 31, 2019

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

DECENTRALIZED SIGNAL PROCESSING AND DISTRIBUTED CONTROL OF COLLABORATIVE AUTONOMOUS NETWORKS

Ryan Goldhahn (17-ERD-101)

IM release #LLNL-TR-795739

Abstract

Intelligent autonomous sensor networks, often comprised of large numbers of sensors, must be capable of jointly exploiting data collected at each agent in the network, and using that data to optimize their future actions towards multiple mission objectives. Centralized signal processing and optimization solutions process all data and determine all future actions at a single agent, and the resulting information and commands are disseminated back to the network. The communications bandwidth this requires and the single point of failure the central agent represents often make these solutions untenable for national security applications. In this project, several fundamental algorithms for solving both the decentralized signal processing and network optimization were developed, as well as simulation software to validate the results of these algorithms at scale. Specifically, novel algorithms for Bayesian decentralized estimation and decentralized detection and optimization based on the alternating direction method of multipliers (ADMM) were developed for autonomous sensor networks and published in the literature. The first large-scale simulation of autonomous sensor networks (1000 agents) was conducted on this project, validating the performance of the developed algorithms. These algorithms and simulation tools are critical components of any decentralized autonomous network and have current and future national security applications, including distributed sensor networks for detection, estimation, and tracking problems, and large decentralized cyber-physical infrastructure such as the power grid.

Background and Research Objectives

In recent years, both sensors and embedded computing systems have become smaller, cheaper, and far more capable. This has enabled smarter sensors capable of better understanding the data they collect, and able to use the information they extract to autonomously determine their future actions, due to the increase in onboard processing. This potential impact of networks of these sensors can be readily seen both in consumer applications such as the the emerging internet-of-things (IoT) and in the national security space. Of particular importance in the national security domain are decentralized solutions, where there is no centralized command and control node or data fusion center. Eliminating this single point of failure from current centralized networks adds resilience to the system but adds significant complexity to the data exploitation and autonomy algorithms. While recent hardware advances have enabled these intelligent networks, the algorithms by which they collectively exploit sensor data and make decisions are still under development. In particular, algorithms are needed which address Byzantine attacks, an important vulnerability of large decentralized networks, where one or more agents in the network is providing false information.

The initial research objectives of this project were thus to develop scalable decentralized algorithms for 1) mapping and estimating a background or environment, 2) detecting of objects of interest and estimating related parameters in the presence of data falsification attacks, 3) optimally repositioning agents for better future performance, and 4) a simulation tool to test and benchmark the performance of the above algorithms at scale. Overall, the project produced novel algorithms towards decentralized detection (Kailkhura et. al 2017), estimation (Ray et. al 2019), and optimization (Schmidt et. al 2019), as well as the first known large-scale simulation of autonomous sensors taking into account real-world communication effects (Yen et. al 2018). However, a ubiquitous issue was the available computation resources; often algorithms could not be developed which would run in real time on low power compute platforms. High fidelity, distributed background mapping in particular was identified after the first year as being infeasible with currently available algorithms and hardware in real time, after a quantitative study on a state-of-the-art low size, weight, and power (SWaP) device (Ho et. al, 2018). Focus was shifted to detection of specific objects or phenomena and the estimation of their related properties. Hardware remained a consideration throughout the project, and some effort was directed at identifying suitable hardware and tailoring algorithms and software to these platforms in the design stage. Preliminary work was also undertaken towards flying these algorithms on multiple unmanned aerial vehicles (UAVs), although issues with hardware and infrastructure limited progress in this area.

Scientific Approach and Accomplishments

This section describes the technical approach, results and achievements towards the above project objectives.

Detection

Detection is generally formulated as a binary hypothesis testing problem, where the observed data is due to measurement noise and background under one hypothesis (H0), and contains some contribution from a target, object, or phenomenon of interest under the other (H1). This is a well-studied problem in the centralized case, in which information is available at a single agent. When parammters of the detection system and signal model are not known, composite hypothesis testing frameworks such as the generalized likelihood ration test (GLRT) are used. However, the GLRT does not have a straightforward decentralized implementation. In decentralized detection approaches, each agent communicates only with its neighbors and updates its local state information about the phenomenon (i.e. a summary statistic) by a local fusion rule that employs a weighted combination of its own value and those received from its neighbors. Agents continue this process until the entire network converges to a steady-state value which is the global test statistic. A simple decentralized target detection solution valid in the low signal-to-noise ratio (SNR) regime was proposed and applied to the problem of detecting a radiation source with unknown location. A decentralized implementation of the derived test which is robust to Byzantine attacks using the alternating direction method of multipliers (ADMM) was derived, and a study of the robustness of the proposed detection algorithm to Byzantine attacks and a comparison with conventional approaches was conducted (Kailkhura et. al 2017). To the best of our knowledge, this was the first such result on Byzantine-resilient locally optimum detection in collaborative autonomous sensor networks.

Thus far, research on detection in the presence of data falsification attacks has primarily focused on the centralized model (Marano et. al 2008). Several attempts have been made to address the security threats in conventional consensus-based detection schemes in recent research (Kailkhura et. al 2016), however the performance analysis of ADMM in the presence of data falsifying Byzantine attacks has thus far not been addressed in the literature. Our work contributed first by rigorously analyzing the effect of erroneous data on the ADMM convergence behavior of multi-agent systems. We showed that the algorithm linearly converges to a neighborhood of the optimal solution under certain conditions and characterized the neighborhood size analytically. We provided guidelines for network design to achieve a faster convergence to the neighborhood. We also provided conditions on the erroneous updates for exact convergence to the optimal solution. Finally, to mitigate the influence of unreliable agents, we proposed a robust decentralized ADMM algorithm (ROAD) and show its resilience to unreliable agents with an exact convergence to the optimum value (Li et. al 2019).

Estimation

In the above detection methods, the steady-state value of a scalar or vector is desired on all agents in a network. In estimation problems, this can correspond to the most likely values of the quantities of interest, given the observed data. Traditional approaches for distributed inference in networks producing only such "point estimates" include factor graphs/sum-product/message passing/belief propagation (Ihler et. al 2005), diffusion (Cattivelli and Sayed 2009), and ADMM (Erseghe 2012). However, it is often advantageous to quantify the uncertainty associated with such an estimate for an understanding of which other values of the parameters of interest may also explain the observed data. A full probability density function (pdf) gives such information and can be used by an intelligent sensor networks to collect future data for better information about the unknown parameters in question. Approximating non-parametric posteriors has traditionally been approached using Markov chain Monte Carlo (MCMC) sampling, which can be inefficient and requires centralized access to all data. A novel method of estimating a full posterior distribution over a large decentralized sensor network has been developed. The method is the only known way to provably compute posterior distributions in a decentralized framework with limited communication between agents, and a record of invention has been recently filed (Ray et. al 2019). Further details are omitted until a patent application has been filed.

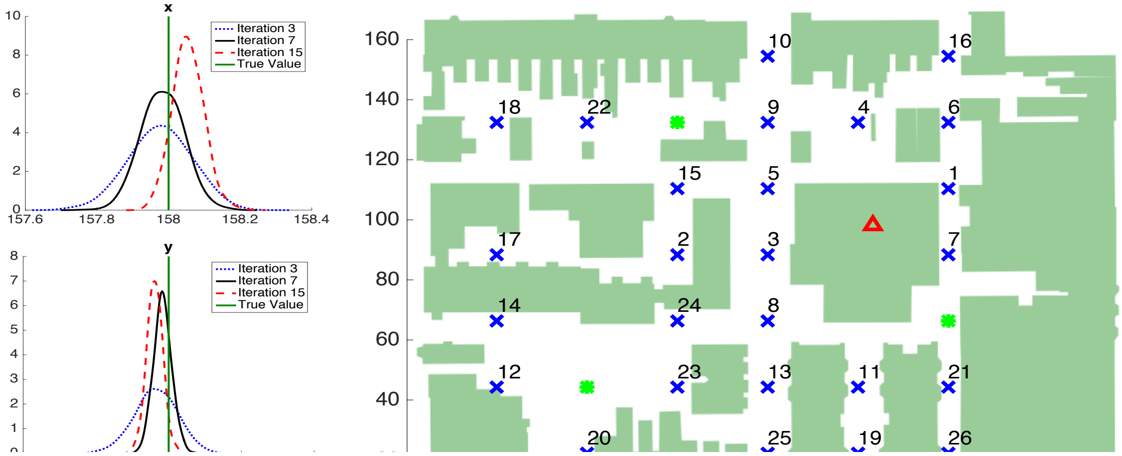

Figure 1: Mobile radiation sensors take measurments to localize a radiation source. The numbered 'x' marks the future positions of radiation sensors sequentially selected to optimally localize a radioactive source in the true position given by the red triangle.

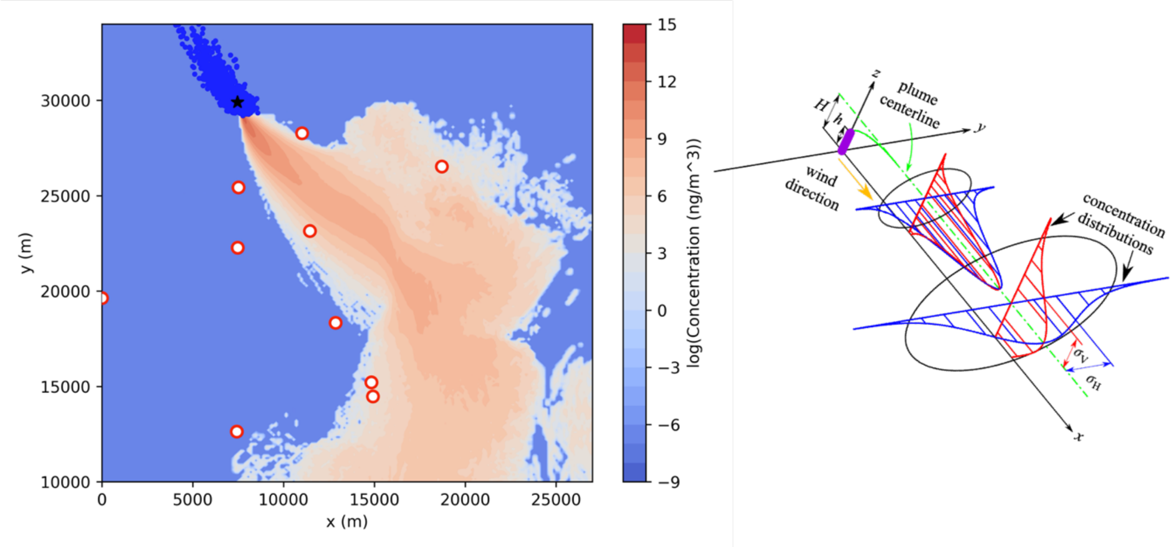

Figure 2: Sensor positions (red circles) are optimally chosen using Shannon entropy to localize the location of the source of a chemical plume using a simple Gaussian plume model to quantify the information gain in future measurements. They select future measurement locations which maximally reduce the uncertainty in the source parameter estimates, shown with blue dots.

Optimization

The above estimation algorithm provides complete information about unknown parameters with a full quantification of uncertainty and limited sharing of data between agents. This is critical for autonomous sensor networks with an estimation problem as an objective; understanding where the uncertainty is located in the parameter space will determine which of a network's future actions will likely produce the most information gain, or the least uncertainty in the estimated parameters, after that data is observed. We proposed a sensing strategy which maximizes mutual information, based on Shannon entropy, to choose the next measurement location from a discrete set of design conditions. Specifically, given a set of experimental measurements $\mathbf{m_{n-1}} = \{\tilde{m}_1, \tilde{m}_2, ..., \tilde{m}_{n-1}\}$, we seek a future action $\xi_n^* \in \Xi$, which would generate new data point $(\xi_n^*, \tilde{m}_n)$, such that we optimally reduce the uncertainty in the parameters when we update the prior distribution using the new set of observations. Let Q denote the random vector of parameters with realizations $q \in \mathcal{Q}$ for p-dimensional parameter space Q. We employ mutual information to choose ξ_n^* from the set of possible future actions Ξ. A utility function based on the conditional Shannon entropy is used to find ξ_n^*:

$$U(m_n, \xi_n) = \int_{\mathcal{Q}} p(q|m_n, \mathbf{m_{n-1}}) \log(p(q|m_n, \mathbf{m_{n-1}}))dq - \int_{\mathcal{Q}} p(q|\mathbf{m_{n-1}}) \log(p(q|\mathbf{m_{n-1}}))dq,$$

which quantifies the reduction in uncertainty provided by the predicted m_n for a yet-to-be-obtained measurement from future sensor action $\xi_n \in \Xi$. We can then compute the average amount of information obtained with action ξ_n^* by marginalizing over the set of all unknown future observations M as

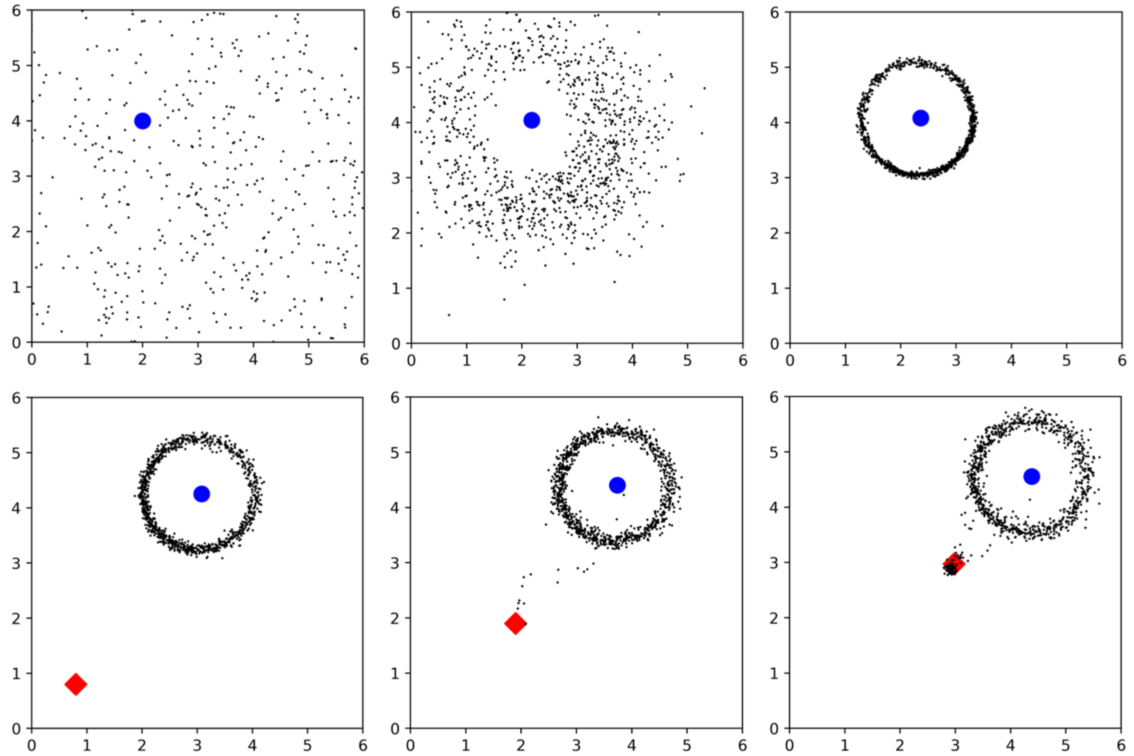

Figure 3: An example of a complex network behavior emerging from a simple objective function and limited information sharing between agents. Agents initially move towards and encircle the blue agent, a subset choose to pursue the red agent.

$$\mathbb{E}_{m_n}[U(m_n,\xi_n)] = \int_{\mathcal{M}} U(m_n,\xi_n)p(m_n|\mathbf{m_{n-1}},\xi_n)dm_n.$$

We apply the movement strategy to two source localization examples: a radiation source in an urban environment (Fig. 1) and a chemical plume release (Fig. 2) (Schmidt et. al 2019). While mutual information is a suitable objective function which is generalizable to many different estimation problems, other mission-driven objectives may also be important to the choice of future actions. Multiple specific objectives may be designed which produce complex behaviors in autonomous sensor networks and optimized using ADMM. The objective function below chooses a future position for an individual sensor using only an estimated target position $\hat{x_t}$, and the positions of the immediate neighbors of the agent in question, $x_{i\in\mathcal{N}(x)}$.

$$F(x) = \underset{\uparrow}{\beta}[a_1 \overbrace{\|\hat{x_t} - x\|_2^2}^{\text{“Chase the red dot”}}] + (1-\underset{\uparrow}{\beta})\left[a_2 \overbrace{\left\|\bar{x}_{i\in\mathcal{N}(x)} - x\right\|_2^2}^{\text{“Keep away from neighbors”}} + a_3 \underbrace{\left(\sum\nolimits_{i\in\mathcal{N}(x)} \|x_i - x\|_2^2 - c\right)^2}_{\text{“Stay distance 1 from the blue dot”}}\right]$$

$$\text{Task assignment} \longrightarrow \beta \sim \text{Bernoulli}(p)$$

The desired behavior included a task assignment problem, with a portion of the agents in the network pursing the red agent, and the remainder staying a unit distance away from a blue agent and avoiding collisions with neighboring agents. The behavior in Fig. 3 is observed, with each agent drawing independently from a Bernoulli distribution with parameter p to decide between these conflicting objectives. While a trivial example, it demonstrates sophisticated network behavior emerging from a simple objective function and requiring only coordination with a small number of nearby agents.

Simulation

While simulation tools are currently available to test single autonomous systems, there is no known capability to simulate large autonomous networks (>1000 agents) under realistic communication conditions. Modelling of communication is particularly important given the degree to which such networks rely on collaboration between agents. Dropped packets and the latency induced by shared, limited bandwidth channels thus substantially affect the performance of any collaborative detection, estimation, or optimization problem. Simulating these effects is critically important when validating the performance of these algorithms at large scales. We developed software based on the ns-3 network simulation software to conduct the first known study of autonomous sensor networks at large scale under realistic communications conditions, and validated the performance of our novel decentralized detection algorithm in a network of 1000 agents, with some of the agents providing false information (Yen et. al 2018).

Additionally, software developed in collaboration with the University of California, San Diego allows the visualization of the high-dimensional data sets collected by large autonomous sensor networks. In Fig. 4, the plume simulation from Fig. 2 is visualized as a point cloud, with both color and particle size a function of the concentration at that point in plume. Individual sensors are shown with a sensor ID ('Sx') and the log of their current measured chemical concentration. Also shown with orange dots is the posterior in latitude and longitude of the source location which could be computed in a decentralized manner using the algorithms described above. The software is fully interactive and allows a user or software developer to query nodes for individual node parameters and data, for better understanding of network behavior, and for high-level control over mission objectives. An intuitive presentation of large amounts of information and the understanding and trust of the end user of the decisions made by the network are absolutely necessary for these networks to be employed in practice.

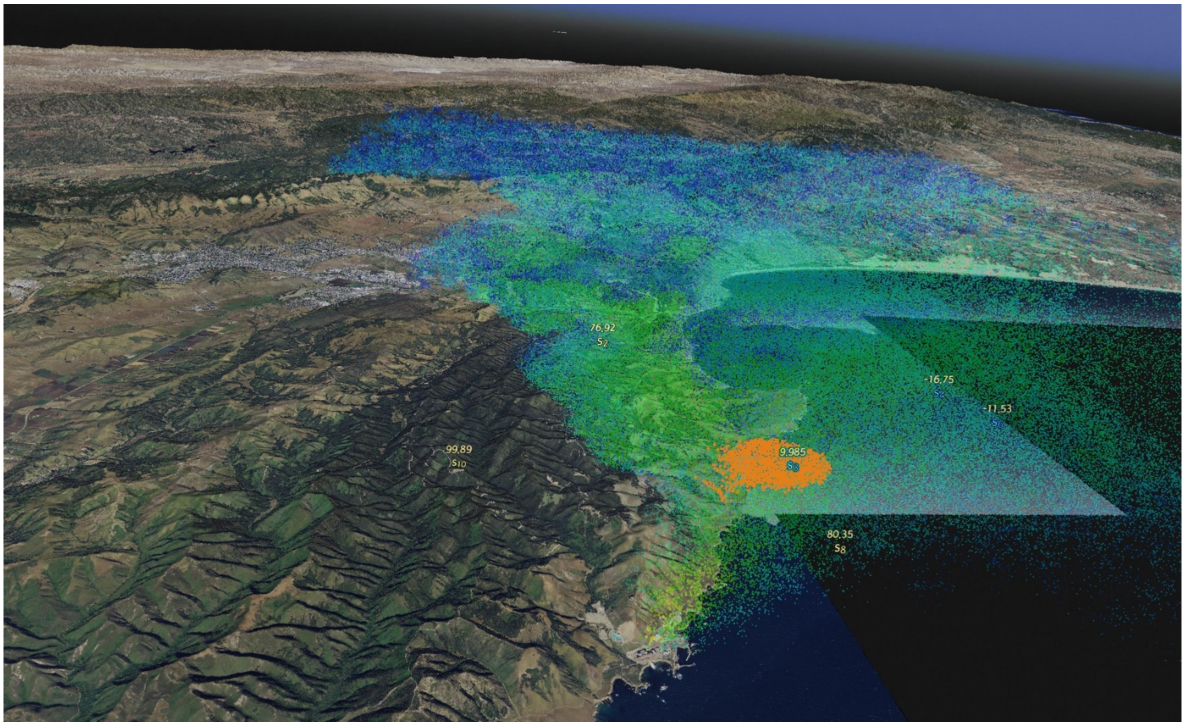

Figure 4: Visualization software was developed for the high dimensional sensor and telemetry data from large autonomous sensor networks. Shown a 3D visualization of the simulation of plume source localization from Fig. 2.

Impact on Mission

This work represents a new research direction for LLNL and has had immediate and direct application to multiple program areas. The algorithms developed under this LDRD are either in use or expected to benefit currently funded programmatic work for the Department of Energy (DOE), National Nuclear Security Administration (NNSA), Department of Homeland Security (DHS), and the Department of Defense (DoD) for decentralized power grid management, detection/estimation, and intelligence, surveillance and reconnaissance (ISR) applications. The success of this project also contributed to the decision to support an institutional investment in a netted enclosure to safely operate unmanned ground and aerial systems. This is an important capability to develop future intelligent autonomous sensor networks, and test, calibrate and benchmark their hardware implementations.

The project also had a significant impact in terms of workforce development, and raised the profile of LLNL in this field. The University of California (UC) San Diego and the University of Texas, Austin were engaged via subcontract to contribute to this work, and summer students were hired from Syracuse University, UC Davis, UC San Diego, and Northeastern University over the course of the project. Two of these summer students were later hired as staff as a direct result, growing LLNL capabilities in embedded systems and decentralized optimization and learning. Professor Hao Chen of Boise State University, a recognized expert in data falsification attacks on decentralized networks, spent a sabbatical year with LLNL because of this work and made important contributions to the project.

Conclusion

This project produced fundamental contributions in multiple important aspects of intelligent

autonomous sensors networks, including 1) the decentralized signal processing allowing agents in the network information from all collected data, 2) the optimization of future actions based on both information theoretic and mission-driven goals, and 3) the simulation tools to validate these algorithms at large scale in real-world conditions.
The means to pursue this work on a theoretical and generalizable basis will be sought, however there is an autonomy component in general, and a collaborative autonomy component in particular, to a large amount of current programmatic work for which autonomy is not the primary focus. Our results may be seen as an enabling capability that will support many critical laboratory mission areas and allow the fielding of cheaper, more robust, and more effective disaggregated solutions. Additional internal and external investment in the general algorithms underlying decentralized data fusion, estimation and network optimization will further improve performance and, in particular, add resilience to more sophisticated adversarial actions.

References

Cattivelli, Federico S., and Ali H. Sayed. "Diffusion LMS strategies for distributed estimation." IEEE Transactions on Signal Processing 58, no. 3 (2009): 1035-1048.

Erseghe, Tomaso. "A distributed and scalable processing method based upon ADMM." IEEE Signal Processing Letters 19, no. 9 (2012): 563-566.

Ho, Nelson, Ryan Goldhahn, Maya Gokhale. "Collaborative Autonomy: Evaluating Feature Extraction on the Nvidia Jetson." LLNL Poster Showcase, 2018.

Hogan, Thomas, Bhavya Kailkhura, Ryan Goldhahn, "Universal Decision-Based Black-Box Perturbations: Breaking Security-Through-Obscurity Defenses," KDD 2019 Workshop on Adversarial Learning Methods for Machine Learning and Data Mining, 2019.

Ihler, Alexander T., John W. Fisher, Randolph L. Moses, and Alan S. Willsky. "Nonparametric belief propagation for self-localization of sensor networks." IEEE Journal on Selected Areas in Communications 23, no. 4 (2005): 809-819.

Kailkhura, Bhavya, Swastik Brahma, and Pramod K. Varshney. "Data falsification attacks on consensus-based detection systems." IEEE Transactions on Signal and Information Processing over Networks 3, no. 1 (2016): 145-158.

Kailkhura, Bhavya, Priyadip Ray, Deepak Rajan, Anton Yen, Peter Barnes, and Ryan Goldhahn. "Byzantine-resilient locally optimum detection using collaborative autonomous networks." In 2017 IEEE 7th international workshop on computational advances in multi-sensor adaptive processing (CAMSAP), pp. 1-5. IEEE, 2017.

Li, Qunwei, Bhavya Kailkhura, Ryan Goldhahn, Priyadip Ray, and Pramod K. Varshney. "Robust decentralized learning using ADMM with unreliable agents." IEEE Transactions on Signal Processing, Submitted (2019). LLNL-CONF-739811-DRAFT

Marano, Stefano, Vincenzo Matta, and Lang Tong. "Distributed detection in the presence of Byzantine attacks." IEEE Transactions on Signal Processing 57, no. 1 (2008): 16-29.

Ray, Priyadip, Hao Chen, Deepak Rajan, Braden Soper, Ryan Goldhahn. "Gradient-based Distributed Bayesian Estimation and Policy Learning in Collaborative Multi-Agent Networks." LLNL Record of Invention, IL-13487, 2019.

Schmidt, Kathleen, Ralph C. Smith, Jason Hite, John Mattingly, Yousry Azmy, Deepak Rajan, and Ryan Goldhahn. "Sequential optimal positioning of mobile sensors using mutual information." Statistical Analysis and Data Mining: The ASA Data Science Journal (2019).

Yen, Anton Y., Peter D. Barnes, Bhavya Kailkhura, Priyadip Ray, Deepak Rajan, Kathleen L. Schmidt, and Ryan A. Goldhahn. "Large-scale parallel simulations of distributed detection algorithms for collaborative autonomous sensor networks." In Disruptive Technologies in Information Sciences, vol. 10652, p. 106520G. International Society for Optics and Photonics, 2018.

Publications and Presentations

Record of Invention

Ray, Priyadip, Hao Chen, Deepak Rajan, Braden Soper, Ryan Goldhahn. "Gradient-based Distributed Bayesian Estimation and Policy Learning in Collaborative Multi-Agent Networks." LLNL Record of Invention, 2019. LLNL-IL-13487

Journal Papers

Schmidt, Kathleen, Ralph C. Smith, Jason Hite, John Mattingly, Yousry Azmy, Deepak Rajan, and Ryan Goldhahn. "Sequential optimal positioning of mobile sensors using mutual information." Statistical Analysis and Data Mining: The ASA Data Science Journal (2019). LLNL-JRNL-753008.

Li, Qunwei, Bhavya Kailkhura, Ryan Goldhahn, Priyadip Ray, and Pramod K. Varshney. "Robust decentralized learning using ADMM with unreliable agents." IEEE Transactions on Signal Processing, Submitted (2019). LLNL-CONF-739811-DRAFT

Conference Papers

Hogan, Thomas, Bhavya Kailkhura, Ryan Goldhahn, "Universal Decision-Based Black-Box Perturbations: Breaking Security-Through-Obscurity Defenses," KDD 2019 Workshop on Adversarial Learning Methods for Machine Learning and Data Mining, 2019. LLNL-CONF-761205.

Kailkhura, Bhavya, Priyadip Ray, Deepak Rajan, Anton Yen, Peter Barnes, and Ryan Goldhahn. "Byzantine-resilient locally optimum detection using collaborative autonomous networks." In 2017 IEEE 7th international workshop on computational advances in multi-sensor adaptive processing (CAMSAP), pp. 1-5. IEEE, 2017. LLNL-CONF-731964.

Yen, Anton Y., Peter D. Barnes, Bhavya Kailkhura, Priyadip Ray, Deepak Rajan, Kathleen L. Schmidt, and Ryan A. Goldhahn. "Large-scale parallel simulations of distributed detection algorithms for collaborative autonomous sensor networks." In Disruptive Technologies in Information Sciences, vol. 10652, p. 106520G. International Society for Optics and Photonics, 2018. LLNL-CONF-749406.

Posters

Ho, Nelson, Ryan Goldhahn, Maya Gokhale. "Collaborative Autonomy: Evaluating Feature Extraction on the Nvidia Jetson." LLNL Poster Showcase, 2018. LLNL-POST-755801.

Kailkhura, Bhavya, Priyadip Ray, Deepak Rajan, Anton Yen, Peter Barnes, Ryan Goldhahn. "Byzantine-Resilient Detection Using Collaborative Autonomous Swarms." IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP), 2017. LLNL-POST-742699.

Schmidt, Katie, Ralph C. Smith, Deepak Rajan, Ryan Goldhahn, Jason Hite, John Mattingly. "Optimal Positioning of Mobile Sensors Using Mutual Information." Conference on Data Analysis (CODA), 2018. LLNL-POST-746706.

Wapman, Jonathan, Priyadip Ray, Bhavya Kailkhura, Ryan Goldhahn. "Chemical Plume Detection with Collaborative Autonomous Sensor Networks." LLNL Signal and Image Sciences (CASIS) Workshop, 2018. LLNL-POST-749008.