AFRL-AFOSR-VA-TR-2019-0027

Real-Time Optimization in Complex Stochastic Environments

Christos Cassandras
TRUSTEES OF BOSTON UNIVERSITY

12/14/2018
Final Report

DISTRIBUTION A: Distribution approved for public release.

Air Force Research Laboratory
AF Office Of Scientific Research (AFOSR)/ RTA2
Arlington, Virginia 22203
Air Force Materiel Command

REPORT DOCUMENTATION PAGE	*Form Approved OMB No. 0704-0188*

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.
PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE *(DD-MM-YYYY)*	**2. REPORT TYPE**	**3. DATES COVERED** *(From - To)*
12/10/2018	Final Report	September 2015 - September 2018

4. TITLE AND SUBTITLE REAL-TIME OPTIMIZATION IN COMPLEX STOCHASTIC ENVIRONMENTS	**5a. CONTRACT NUMBER**
	5b. GRANT NUMBER FA9550-15-1-0471
	5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S) Christos G. Cassandras	**5d. PROJECT NUMBER**
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Boston University, Boston, MA 02215	**8. PERFORMING ORGANIZATION REPORT NUMBER** N/A
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Optimization and Discrete Mathematics AFOSR 875 North Randolph Street Suite 325 Arlington, VA 22203	**10. SPONSOR/MONITOR'S ACRONYM(S)**
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Public availability

13. SUPPLEMENTARY NOTES

14. ABSTRACT
The research in this project aims to enable a systematic on-line use of optimization techniques for real-time applications in complex, time-varying stochastic environments. This is in contrast to off-line optimization where computationally intensive methods may be used. The main outcomes of the project are: (a) event-driven distributed optimization algorithms which exploit the spatial decomposition of complex optimization problems and involve minimal communication, (b) receding horizon algorithms which exploit a temporal (rather than spatial) decomposition, (c) a novel "boosting function" approach which takes advantage of structural information to escape local optima and approach global optimality.

15. SUBJECT TERMS
Distributed Optimization Algorithms, Asynchronous Optimization Algorithms, Event-driven Optimization, Multi-Agent Systems, Receding Horizon Algorithms, Perturbation Analysis, Stochastic Hybrid Systems

16. SECURITY CLASSIFICATION OF:			**17. LIMITATION OF ABSTRACT**	**18. NUMBER OF PAGES**	**19a. NAME OF RESPONSIBLE PERSON**
a. REPORT	**b. ABSTRACT**	**c. THIS PAGE**			Christos G. Cassandras
Unclassified	Unclassified	Unclassified	UU	15	**19b. TELEPHONE NUMBER** *(Include area code)* 617-353-7154

INSTRUCTIONS FOR COMPLETING SF 298

1. REPORT DATE. Full publication date, including day, month, if available. Must cite at least the year and be Year 2000 compliant, e.g. 30-06-1998; xx-06-1998; xx-xx-1998.

2. REPORT TYPE. State the type of report, such as final, technical, interim, memorandum, master's thesis, progress, quarterly, research, special, group study, etc.

3. DATE COVERED. Indicate the time during which the work was performed and the report was written, e.g., Jun 1997 - Jun 1998; 1-10 Jun 1996; May - Nov 1998; Nov 1998.

4. TITLE. Enter title and subtitle with volume number and part number, if applicable. On classified documents, enter the title classification in parentheses.

5a. CONTRACT NUMBER. Enter all contract numbers as they appear in the report, e.g. F33315-86-C-5169.

5b. GRANT NUMBER. Enter all grant numbers as they appear in the report. e.g. AFOSR-82-1234.

5c. PROGRAM ELEMENT NUMBER. Enter all program element numbers as they appear in the report, e.g. 61101A.

5e. TASK NUMBER. Enter all task numbers as they appear in the report, e.g. 05; RF0330201; T4112.

5f. WORK UNIT NUMBER. Enter all work unit numbers as they appear in the report, e.g. 001; AFAPL30480105.

6. AUTHOR(S). Enter name(s) of person(s) responsible for writing the report, performing the research, or credited with the content of the report. The form of entry is the last name, first name, middle initial, and additional qualifiers separated by commas, e.g. Smith, Richard, J, Jr.

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES). Self-explanatory.

8. PERFORMING ORGANIZATION REPORT NUMBER. Enter all unique alphanumeric report numbers assigned by the performing organization, e.g. BRL-1234; AFWL-TR-85-4017-Vol-21-PT-2.

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES). Enter the name and address of the organization(s) financially responsible for and monitoring the work.

10. SPONSOR/MONITOR'S ACRONYM(S). Enter, if available, e.g. BRL, ARDEC, NADC.

11. SPONSOR/MONITOR'S REPORT NUMBER(S). Enter report number as assigned by the sponsoring/ monitoring agency, if available, e.g. BRL-TR-829; -215.

12. DISTRIBUTION/AVAILABILITY STATEMENT. Use agency-mandated availability statements to indicate the public availability or distribution limitations of the report. If additional limitations/ restrictions or special markings are indicated, follow agency authorization procedures, e.g. RD/FRD, PROPIN, ITAR, etc. Include copyright information.

13. SUPPLEMENTARY NOTES. Enter information not included elsewhere such as: prepared in cooperation with; translation of; report supersedes; old edition number, etc.

14. ABSTRACT. A brief (approximately 200 words) factual summary of the most significant information.

15. SUBJECT TERMS. Key words or phrases identifying major concepts in the report.

16. SECURITY CLASSIFICATION. Enter security classification in accordance with security classification regulations, e.g. U, C, S, etc. If this form contains classified information, stamp classification level on the top and bottom of this page.

17. LIMITATION OF ABSTRACT. This block must be completed to assign a distribution limitation to the abstract. Enter UU (Unclassified Unlimited) or SAR (Same as Report). An entry in this block is necessary if the abstract is to be limited.

FINAL REPORT

AFOSR GRANT: FA9550-15-1-0471

TITLE: **REAL-TIME OPTIMIZATION IN COMPLEX STOCHASTIC ENVIRONMENTS**

REPORTING PERIOD: September 2015 – September 2018

Principal Investigator:

Christos G. Cassandras
Division of Systems Engineering
Boston University
Boston, MA 02215
Tel: (617) 353-7154
FAX: (617) 353-4830
e-mail: *cgc@bu.edu*

December 2018

1. OBJECTIVES

The research pursued under grant FA9550-15-1-0471 aims at enabling a systematic *on-line* use of optimization techniques (as opposed to off-line methods which can afford to be computationally intensive) for real-time applications in complex stochastic environments. Our work has led to a framework of novel complementary methods addressing this limitation while recognizing requirements for new generations of systems critical to the national infrastucture and consistent with emerging data-driven, network-centric views of warfare. This framework rests on three conceptual cornerstones: the event-driven paradigm for optimization, data-driven methodologies for optimization algorithms, and scalable on-line solutions for optimal control problems.

Four major research objectives were pursued:

1. Event-driven distributed optimization based on spatial decomposition. For settings where agents form a team cooperating towards a common objective, it is natural to model the agents as nodes in a network and to formulate optimization problems subject to constraints that pertain to the capabilities of these nodes (e.g., mobility, sensing range). Depending on the network configuration, a spatial decomposition is possible by associating each node with a *neighborhood* of other nodes and limiting the exchange of information to nodes communicating only with their respective neighborhoods. This exchange of information is traditionally based on a synchronization scheme such that nodes periodically communicate. However, this is inefficient and, in fact, unnecessary since often the state of a node may not have changed or may have only changed in a predictable way. Our approach, therefore, has been to develop asynchronous (specifically, *event-driven*) approaches by controlling the exchange of information among nodes in a networked environment to occur only when absolutely necessary, dictated by "events" defined through certain well-defined conditions.

2. **Event-driven receding-horizon optimization based on temporal decomposition.** In contrast to a spatial decomposition, a temporal decomposition approach (suitable for uncertain environments where real-time solutions of combinatorially complex assignment and routing algorithms are infeasible) is based on the idea of breaking down the solution of a complex stochastic problem into a number of smaller deterministic problems over time. Each problem is solved only when new data are obtained or new random events are observed. In this way, our approach to *receding horizon* optimization, unlike conventional receding horizon methods which require continuous time-driven iterations, is to update a control action only when certain events occur. This can drastically reduce computation with minimal effect on optimality.

3. **Escaping local optima in real-time optimization**. The distributed and receding horizon optimization methods mentioned above are based on objective function gradients either computed or estimated (in stochastic settings). The objective functions in the problems we are targeting are almost always nonlinear and nonconvex. This implies the existence of multiple local optima. Consequently, gradient-based methods are limited by their inability to escape such points and determine global optima. In this project, we have pursued a new research direction based on a "boosting function" concept which takes advantage of some structural information regarding the state space of a given problem. In a network setting, the main idea is to temporarily alter the objective function of a node whenever an equilibrium is reached. A boosting function is a transformation of the associated local partial derivative (which is zero at an equilibrium point) to a new derivative which is non-zero and, therefore, forces a node to move in a direction determined by the boosting function and to explore the search space for a an improved equilibrium (which is potentially, but not necessarily, a global optimum).

4. **On-line solutions for optimal control problems**. This research objective goes beyond the parametric optimization problems considered thus far. The goal here is to formulate *dynamic* optimization problems, where control variables explicitly affect the dynamics of agents or the environment itself is subject to its own dynamics. Optimal control formulations are ideal for such a setting, except for their prohibitive computational complexity, especially when used on line. In this project, we have shown how to bypass this limitation by exploiting the fact that the structure of an optimal control problem solution often has a parametric form. When it does not, one can still define parametric families of solutions which can be optimized and yield near-optimal or at least vastly improved solutions relative to ad hoc policies often adopted.

2. ACCOMPLISHMENTS AND NEW FINDINGS

2.1. Event-driven Optimization of Multi-Agent Systems: the Event Excitation Problem.

Event-driven optimization schemes in a networked multi-agent environment have the advantage of requiring only limited interactions between the agents, based on specific events which trigger significant enough changes to warrant inter-node communication. Clearly, the premise of these methods is that the events involved are observable so as to "excite" an underlying event-driven controller. However, it is not always obvious that these events actually take place under every feasible control: it is possible that under some control no such events are excited, in which case the controller may be useless. In such cases, one can resort to artificial "timeout events" so as to eventually take actions, but this is obviously inefficient. Moreover, in event-driven optimization mechanisms this problem results in very slow convergence to an optimum or in an algorithm failing to generate any improvement in the decision variables being updated.

While developing event-driven optimization algorithms for some specific multi-agent system problems of interest, we discovered this particular issue by observing that gradients of our objective functions would frequently remain at or close to zero values. This obviously limits the use of event-driven gradient-based algorithms. To address this problem, we created a new metric for the objective function which creates a potential field guaranteeing non-zero gradient values when no events are present, as well as eventual event excitation [2]. Once this event excitation problem was resolved, we were able to successfully develop event-driven gradient-based optimization algorithms for a large class of problems. An example is the data harvesting problem [34], where N mobile agents are tasked with the collection of data generated at M stationary sources and delivery to a base with the goal of minimizing expected collection and delivery delays.

2.2. Event-Driven Cooperative Receding Horizon Optimization for Multi-agent Systems in Uncertain Environments.

In considering multi-agent systems operating in uncertain environments, we need to model the unknown future. This is typically accomplished with explicit stochastic models or through simulation. The former approach is challenging because we must often make questionable assumptions, lack the necessary information, or resort to unreliable data to construct appropriate probability distributions. Simulation, on the other hand, is a time consuming task. Both are "estimate-and-plan" approaches in which a policy is derived to dictate optimal decisions to make in real time as a function of an observed system state. An alternative is a "hedge-and-react" approach whereby we simply wait for a random event to occur and then react, provided that the decision making process upon reaction is sufficiently fast relative to the random event frequency. This amounts to a decomposition of a complex problem over time: we solve an optimization problem over a given time horizon, and then continuously extend this time horizon forward. This is the basis of the Receding Horizon (RH) optimization framework which we have developed in our prior work. In the course of this project, we addressed a number of limitations of our original framework and developed a new event-driven Cooperative Receding Horizon (CRH) approach [27] to solve maximum reward collection problems where multiple agents cooperate to maximize the total reward collected from a set of targets in a given mission space. In this new CRH framework a controller sequentially solves optimization problems over a planning horizon and executes the control for a shorter action horizon, where both are defined by certain events associated with new information becoming available. One of the key accomplishments was a reduction of the computational complexity involved in each of the optimization problems, which renders the approach much more amenable to real-time implementation.

2.3. Event-Driven Algorithms for Dynamic Optimization in Cooperative Multi-Agent Systems.

The optimization problems considered above are all *static* or *parametric*, rather than *dynamic*. We have found that developing distributed algorithms for such dynamic problems is much more challenging. A case in point arises in the optimal dynamic formation problem we solved in [9]. This problem arises in mobile leader–follower networks where an optimal formation is generated to maximize a given objective function while continuously preserving connectivity. For the class of optimal formation problems where the objective is to maximize the coverage of a given region of interest, we have shown that the optimal formation is a tree which can be efficiently constructed without solving what in general is a mixed integer nonlinear problem (MINLP). Even so, however, it remains difficult to develop distributed algorithms for such optimal formation problems. On the other hand, treating Connected Automated Vehicles (CAVs) as agents in a transportation systems, we were able to obtain solutions to optimal traffic control problems which are distributed and potentially event-driven [30],[24].

Equally challenging to the optimal formation problem in [9] is the class of persistent monitoring dynamic optimization problems in which the movement of agents is continuously controlled so that, as a cooperative team, they can minimize an uncertainty metric associated with a finite number of data sources. In a one-dimensional mission space, we have shown [39] that the solution can be reduced to a simpler parametric optimization problem: determining a sequence of locations where each agent may dwell for a finite amount of time and then switch direction. Using the infinitesimal perturbation analysis (IPA) methodology for gradient estimation, we were able to obtain a complete on-line solution through an event-driven gradient-based algorithm making use of the new metric to deal with event excitation problem discussed in Section 2.1. We were also able to extend this event-driven optimization scheme to a two-dimensional mission space where agents are restricted to move on a graph topology consisting of multiple intersecting line segments (subspaces) on such a graph [17], as well as to the case where the agents have second-order dynamics [44].

Towards the end of the project, we were able to derive "almost decentralized" algorithms for the class of persistent monitoring problems and, more generally, all multi-agent optimization problems where an agent's trajectory is designed in terms of general function families characterized by parameters that we can optimize. In particular, we first identified conditions under which the *centralized* solution to the optimal multi-agent persistent monitoring problem can be recovered in a *decentralized* event-driven manner. The key to these conditions rests on the fact that the solution to the centralized problem (first obtained in [39]) depends on an event-driven gradient-based algorithm making use of IPA. We can show that the IPA gradient can be recovered in a distributed manner based on local information, except for one event requiring communication from a non-neighbor agent. This has opened up the possibility of developing an entire class of "almost decentralized" algorithms. Alternatively, in our recent work in [53], rather than parameterizing agent trajectories, we adopt a point of view whereby targets are nodes in a graph and their connectivity defines feasible paths in the graph. An agent trajectory in such a graph is specified by a sequence of nodes and an associated dwell time at each node in the sequence. By defining threshold parameters on these dwell times, we have been able to define a class of event-driven controllers which is distributed and scalable. By adjusting the thresholds, we can control the agent behavior in terms of target visiting and dwelling and, therefore, optimize agent performance within the specific parametric controller family considered.

Based on the most recent results summarized above, we believe to be at the cusp of developing a new class of distributed or "almost distributed" optimization algorithms for networked multi-agent systems.

2.4. Escaping Local Optima in Distributed Optimization.

All along this project, we have found that many of the most interesting problems we need to address are non-convex and exhibit multiple local optima. The large class of coverage control problems arising in multi-agent cooperative systems is a case in point. One way we have attempted to deal with this problem is by exploiting the submodularity properties of many objective functions of interest and deriving improved bounds which bring us very close (sometimes within 5%) to the global optimum through the use of simple greedy algorithms as shown in our recent work [40].

In a more general setting (where submodularity is not present), we have pursued a systematic approach for escaping a local optimum, rather than randomly perturbing controllable variables away from it. By exploiting the structure of a given problem, we define "boosting functions" which transform a node's local partial derivative at an equilibrium point (where its value is zero) into a new nonzero value which induces nodes to explore poorly covered areas of the mission space until a new equilibrium point is reached. This boosting process ensures that, at its conclusion, the objective function is no worse than its pre-boosting value (even though a global optimum can still not be guaranteed). We have also conducted extensive simulation-based experiments to demonstrate how this approach improves the solutions obtained through a variety of optimization problems. However, it is only near the end of this project that we were able to make two major advances: (a) Rigorously formalizing the overall boosting process, and, more importantly, (b) Showing that a *distributed* optimization process based on boosting functions is possible in a very general setting. Thus, we have established the foundations for an Asynchronous Distributed Boosting (ADB) optimization framework which can provably converge and result in improved solutions when multiple local optima exist. In the course of this work, we also developed a new interactive simulation environment, available at http://www.bu.edu/codes/research/distributed-control/ which includes ADB optimization. We also continue to use a laboratory test bed with small wireless mobile robots acting as "agents" in a "real-world" cooperative multi-agent setting (see http://www.bu.edu/codes/platforms/.)

3. PUBLICATIONS RESULTED FROM FA9550-15-1-0471

• Papers Published:

[1] Brisimi, T.S., Cassandras, C.G., Osgood, C., Paschalidis, I.C., and Zhang, Y., "Sensing and Classifying Roadway Obstacles in Smart Cities: The *Street Bump* System", *IEEE Access*, 4, pp. 1301-1312, 2016.

[2] Khazaeni, Y., and Cassandras, C.G., "Event excitation for event-driven control and optimization of multi-agent systems", *Proc. of 2016 Intl. Workshop on Discrete Event Systems*, pp. 197-202, 2016.

[3] Mao, J., and Cassandras, C.G., "Solving A Class of Discrete Event Simulation-based Optimization Problems Using "Optimality in Probability"", *Proc. of 2016 Intl. Workshop on Discrete Event Systems*, pp. 129-134, 2016.

[4] Zhang, Y., Malikopoulos A.A., and Cassandras, C.G., "Optimal Control and Coordination of Connected and Automated Vehicles at Urban Traffic Intersections", *Proc. of 2016 American Control Conference,* pp. 6227-6232, 2016.

[5] Fleck, J.L., and Cassandras, C.G., "Personalized Cancer Therapy Design: Robustness vs. Optimality", *Proc. of 55th IEEE Conference on Decision and Control*, pp. 5041-5046, 2016.

[6] Zhou, N., Yu, X., Andersson, S.B., and Cassandras, C.G., "Optimal Event-Driven Multi-Agent Persistent Monitoring of a Finite Set of Targets", *Proc. of 55th IEEE Conference on Decision and Control*, pp. 1814-1819, 2016.

[7] Zhang, J., Pourazarm, S., Cassandras, C.G., and Paschalidis, I.C., "The Price of Anarchy in Transportation Networks by Estimating User Cost Functions from Actual Traffic Data", *Proc. of 55th IEEE Conference on Decision and Control*, pp. 789-794, 2016.

[8] Cassandras, C.G., "On-line Control and Optimization for Cyber-Physical Systems", in *Cyber-Physical Systems: Foundations, Principles and Applications* (H. Song, D.B. Rawat, S. Jeschke, C. Brecher, Eds.), pp. 31-54, Elsevier, 2016.

[9] Sun, X., and Cassandras, C.G., "Optimal Dynamic Formation Control of Multi-Agent Systems in Constrained Environments", *Automatica*, Vol. 73, No. 11, pp. 169-179, 2016.

[10] Pourazarm, S., and Cassandras, C.G., "Energy-based Lifetime Maximization and Security of Wireless Sensor Networks with General Non-ideal Battery Models", *IEEE Trans. on Control of Network Systems*, Vol. 4, No. 2, pp. 323-335, 2017.

[11] Fleck, J.L., and Cassandras, C.G., “Optimal Design of Personalized Prostate Cancer Therapy using Infinitesimal Perturbation Analysis”, *Nonlinear Analysis: Hybrid Systems*, Vol. 25, pp. 246-262, 2017.

[12] Zhang, Y., Cassandras, C.G., and Malikopoulos A.A., “Optimal Control of Connected Automated Vehicles at Urban Traffic Intersections: A Feasibility Enforcement Analysis”, *Proc. of 2017 American Control Conference*, pp. 3548-3553, 2017.

[13] Yu, X., Andersson S.B., Zhou, N. and Cassandras, C.G., “Optimal Dwell Times for Persistent Monitoring of a Finite Set of Targets”, *Proc. of 2017 American Control Conference*, pp. 5544-5549, 2017.

[14] Cassandras, C.G., “Complexity Made Simple (at a Small Price)”, *Proc. of 38th Intl. Conference Petri Nets 2017*, pp. 12-18, 2017.

[15] Zhang, J., Pourazarm, S., Cassandras, C.G., and Paschalidis, I.C., “Data-driven Estimation of Origin-Destination Demand and User Cost Functions for the Optimization of Transportation Networks”, *Proc. of 20th IFAC World Congress*, pp. 10090-10095, 2017.

[16] Wardi, Y. Cassandras, C.G., and Cao, X.R., “Perturbation Analysis: A Framework for Data-Driven Control and Optimization of Discrete Event and Hybrid Systems”, *Proc. of 20th IFAC World Congress*, pp. 3083-3093, 2017.

[17] Zhou, N., Cassandras, C.G., Yu, X., and Andersson S.B., “Optimal Event-Driven Multi-Agent Persistent Monitoring with Graph-Limited Mobility”, *Proc. of 20th IFAC World Congress*, pp. 2217-2222, 2017.

[18] Cassandras, C.G., “Automating Mobility in Smart Cities”, *Proc. of 20th IFAC World Congress*, pp. 7194-7197, 2017.

[19] Swaszek, R.M.A., and Cassandras, C.G., “A Receding Horizon Routing Controller for Inventory Replenishment Agents in Bike Sharing Systems with Time-Varying Demand”, *Proc. of 2017 IEEE Conference on Control Technology and Applications*, pp. 1587-1592, 2017.

[20] Cassandras C.G., “Automating Mobility in Smart Cities”, *Annual Reviews in Control*, Vol. 44, pp. 1-8, 2017.

[21] Pourazarm, S., and Cassandras C.G., “Optimal Routing for Lifetime Maximization of Wireless Sensor Networks with a Mobile Source Node”, *IEEE Trans. on Control of Network Systems*, Vol. 4, No. 4, pp. 793-804, 2017.

[22] Sun, X., Cassandras, C.G., and Meng, X., "A Submodularity-Based Approach for Multi-Agent Optimal Coverage Problems", *Proc. of 56th IEEE Conference on Decision and Control*, pp. 4082-4087, 2017.

[23] Zhou, N., Cassandras, C.G., Yu, X., and Andersson S.B., "Decentralized Event-Driven Algorithms for Multi-Agent Persistent Monitoring Tasks", *Proc. of 56th IEEE Conference on Decision and Control*, pp. 4064-4069, 2017.

[24] Zhang, Y., Malikopoulos A.A., and Cassandras, C.G., "Decentralized Optimal Control for Connected Automated Vehicles at Intersections Including Left and Right Turns", *Proc. of 56th IEEE Conference on Decision and Control*, pp. 4428-4433, 2017.

[25] Chen, R. and Cassandras, C.G., "Stochastic Flow Models with Delays and Applications to Multi-Intersection Traffic Light Control", *Proc. of 2018 Intl. Workshop on Discrete Event Systems*, pp. 50-55, 2018.

[26] Zhang, Y., Cassandras, C.G., Li, W., and Mosterman, P., "A Discrete-Event and Hybrid Simulation Framework Based on SimEvents for Intelligent Transportation System Analysis", *Proc. of 2018 Intl. Workshop on Discrete Event Systems*, pp. 334-339, 2018.

[27] Khazaeni, Y., and Cassandras C.G., "Event-Driven Cooperative Receding Horizon Control for Multi-agent Systems in Uncertain Environments", *IEEE Trans. on Control of Network Systems*, Vol. 5, No. 1, pp. 409-422, 2018.

[28] Mao, J., and Cassandras, C.G., "Solving A Class of Discrete Event Simulation-based Optimization Problems Using "Optimality in Probability"", *J. of Discrete Event Dynamic Systems*, Vol. 28, 1, pp. 35-61, 2018.

[29] Zhang, J., Pourazarm, S., Cassandras C.G., and Paschalidis, I. Ch., "The Price of Anarchy in Transportation Networks: Data-Driven Evaluation and Reduction Strategies", Vol. 106, 4, *Proceedings of the IEEE*, pp. 538-553, 2018.

[30] Malikopoulos, A.A., Cassandras C.G., and Zhang, Y.J., "A Decentralized Energy-Optimal Control Framework for Connected Automated Vehicles at Signal-Free Intersections", *Automatica*, Vol. 93, 7, pp. 244-256, 2018.

[31] Wardi, Y. Cassandras, C.G., and Cao, X.R., "Perturbation Analysis: A Framework for Data-Driven Control and Optimization of Discrete Event and Hybrid Systems", *Annual Reviews in Control*, Vol. 45, pp. 267-280, 2018.

[32] Nenchev, V., Cassandras, C.G., and Raisch, J., "Optimal Control for a Robotic Exploration, Pick-up and Delivery Problem", *Nonlinear Analysis: Hybrid Systems*, Vol. 30, pp. 266-284, 2018.

[33] Pourazarm, S., and Cassandras C.G., "Optimal Routing of Energy-aware Vehicles in Transportation Networks with Inhomogeneous Charging Nodes", *IEEE Trans. on Intelligent Transportation Systems*, Vol. 18, 8, pp. 2515-2527, 2018.

[34] Khazaeni, Y., and Cassandras C.G., "Event-driven Trajectory Optimization for Data Harvesting in Multi-Agent Systems", *IEEE Trans. on Control of Network Systems*, Vol. 5, No. 3, pp. 1335-1348, 2018.

[35] Yu, X., Andersson S.B., Zhou, N. and Cassandras, C.G., "Optimal visiting schedule search for persistent monitoring of a finite set of targets", *Proc. of 2018 American Control Conference*, pp. 4032-4037, 2018.

[36] Meng, X., Houshmand, A., and Cassandras, C.G., "Hybrid System Modeling of Multi-Agent Coverage Problems with Energy Depletion and Repletion", *Proc. of 6th IFAC Conference on Analysis and Design of Hybrid Systems*, pp. 223-228, 2018.

[37] Zhang, Y., and Cassandras, C.G., "The Penetration Effect of Connected Automated Vehicles in Urban Traffic: an Energy Impact Study", *Proc. of 2018 IEEE Conference on Control Technology and Applications*, pp. 620-625, 2018.

[38] Houshmand, A., and Cassandras, C.G., "Eco-Routing of Plug-In Hybrid Electric Vehicles in Transportation Networks", *Proc. of 21st IEEE Intl. Conference on Intelligent Transportation Systems*, 2018.

[39] Zhou, N., Yu, X., Andersson, S.B., and Cassandras C.G., "Optimal Event-Driven Multi-Agent Persistent Monitoring of a Finite Set of Data Sources", *IEEE Trans. on Automatic Control*, Vol. 63, 12, pp. 4204-4217, 2018.

[40] Sun, X., Cassandras C.G., and Meng, X., "Exploiting Submodularity to Quantify Near-Optimality in Multi-Agent Coverage Problems", *Automatica*, Vol. 100, 2, pp. 349-359, 2019.

• Accepted, but not yet published:

[41] Zhang, Y., and Cassandras, C.G., "A Decentralized Optimal Control Framework for Connected Automated Vehicles at Urban Intersections with Dynamic Resequencing", to appear in *Proc. of 57th IEEE Conference on Decision and Control*, 2018.

[42] Meng, X., and Cassandras, C.G., "Optimal Control of Autonomous Vehicles for Non-Stop Signalized Intersection Crossing", to appear in *Proc. of 57th IEEE Conference on Decision and Control*, 2018.

[43] Meng, X., Houshmand, A., and Cassandras, C.G., "Multi-Agent Coverage Control with Energy Depletion and Repletion", to appear in *Proc. of 57th IEEE Conference on Decision and Control*, 2018.

[44] Wang, Y-W., Wei, Y-W., Liu, X-K., Zhou, N., and Cassandras C.G., "Optimal Persistent Monitoring Using Second-Order Agents with Physical Constraints", to appear in *IEEE Trans. on Automatic Control*, 2019.

• Submitted, but not yet accepted:

[45] Yu, X., Andersson, S.B., Zhou, N., and Cassandras C.G., "Scheduling Multiple Agents in a Persistent Monitoring Task Using Reachability Analysis", subm. to *IEEE Trans. on Automatic Control*, 2018.

[46] Feng, Z., Hu, G., and Cassandras, C.G., "Finite-Time Distributed Convex Optimization Under Uncertain Information: A Robust Consensus Approach", subm. to *IEEE Trans. on Control of Network Systems*, 2018.

[47] Meng, X., and Cassandras, C.G., "Eco-driving of Autonomous Vehicles for Non-stop Crossing of Signalized Intersections", subm. to *Transportation Research Part C*, 2018.

[48] Zhang, Y., and Cassandras, C.G., "An Impact Study of Integrating Connected Automated Vehicles with Conventional Traffic", subm. to *Annual Reviews in Control*, 2018.

[49] Chen, R., and Cassandras, C.G., "Stochastic Flow Models with Delays, Blocking and Applications to Multi-Intersection Traffic Light Control", subm. to *J. of Discrete Event Dynamic Systems*, 2018.

[50] Zhang, Y., Cassandras, C.G., Li, W., and Mosterman, P.J., "A Discrete-Event and Hybrid Traffic Simulation Model Based on SimEvents for Intelligent Transportation System Analysis in Mcity", subm. to *J. of Discrete Event Dynamic Systems*, 2018.

[51] Meng, X., and Cassandras, C.G., "A Real-Time Optimal Eco-driving for Autonomous Vehicles Crossing Multiple Signalized Intersections", subm. to *2019 American Control Conference*, 2018.

[52] Xiao, W., and Cassandras, C.G., "Decentralized Optimal Merging Control for Connected and Automated Vehicles", subm. to *2019 American Control Conference*, 2018.

[53] Zhou, N., Cassandras, C.G., Yu, X., and Andersson S.B., "Optimal Threshold-Based Distributed Control Policies for Persistent Monitoring on Graphs", subm. to *2019 American Control Conference*, 2018.

[54] Xiao, W., Belta, C. and Cassandras, C.G., "Decentralized Optimal Merging at an Intersection: a Control Barrier Function Approach", subm. to *10th ACM/IEEE Intl. Conference on Cyber-Physical Systems*, 2018.

4. PERSONNEL SUPPORTED

• Principal Investigator:

Christos G. Cassandras, Professor, Boston University

• Graduate Students:

- Yasaman Khazaeni (PhD obtained, 2016)
- Julia Lima Fleck (PhD obtained, 2016)
- Sepideh Pourazarm (PhD obtained, 2017)
- Xinmiao Sun (PhD obtained, 2017)
- Yue Zhang (expected to complete, 2018)
- Nan Zhou (expected to complete, 2018)

The PhD dissertation completed by **Yasaman Khazaeni** is entitled "An Event-Driven Approach to Control and Optimization of Multi-Agent Systems". It studies event-driven control schemes for multi-agent systems. A new cooperative receding horizon (CRH) controller is designed and applied to a class of maximum reward collection problems where target rewards are time-variant with finite deadlines. The environment is uncertain in the sense that a target's arrival time is not fully known and the agents may have a finite sensing range to detect targets. Unlike traditional time-driven receding horizon control methods, the new methodology adopts an event-driven approach to updating control actions. Controls are optimized for a planning horizon and executed for a shorter action horizon. The dynamic action horizon is determined given the new events during the mission. Assuming no prior information about the uncertainties of the environment, the controller considers each piece of information as a new event. Several limitations of a previously developed CRH controller are addressed resulting in improved computational cost of the controller. The data harvesting problem is also studied in the context of multi-agent systems. Here a set of agents collect data from a finite number of targets and deliver them to the base. First a rigorous optimal control solution is employed using numerical solutions which turn out to be computationally infeasible in real time applications. Using the structures of the optimal control solution, parametric families of trajectories which can be optimized using gradient-based techniques are studied. In order to estimate the gradient of the underlying performance measure, Infinitesimal Perturbation Analysis is utilized which relies on the event-driven nature of the underlying system. The premise of event-driven methods is that the events involved are observable so as to "excite" the underlying event-driven controller. However, it is not always obvious that these events actually take place under every feasible control, in which case the controller may be useless. The issue of event excitation which arises in problems such as data harvesting is studied and addressed by introducing a novel performance measure which generates a potential field over the mission space. The effectiveness of this method has been demonstrated with several examples using different parametric families for the agent trajectories.

The PhD dissertation completed by **Julia Lima Fleck** is entitled "Topics in Perturbation Analysis for Stochastic Hybrid Systems". It studies optimization problems arising in Stochastic Hybrid Systems (SHS) where the size and complexity of SHS frequently render the use of exhaustive verification techniques prohibitive. In this context, Perturbation Analysis techniques, and in particular Infinitesimal Perturbation Analysis

(IPA), have proven to be particularly useful for this class of systems. They are applied in the thesis to two different problems: Traffic Light Control (TLC) and control of cancer progression, both of which are viewed as dynamic optimization problems in an SHS environment. For the TLC problem, a quasi-dynamic control policy is developed based on partial state information defined by detecting whether vehicle backlogs are above or below certain controllable threshold values. Online gradient estimates of a cost metric with respect to these controllable parameters are derived using IPA techniques. These estimators are subsequently used to iteratively adjust the threshold values so as to improve overall system performance. This quasi-dynamic analysis is subsequently extended to parameterize the control policy by green and red cycle lengths as well as queue content thresholds. In the second part of the thesis, the problem of controlling cancer progression is formulated within a Stochastic Hybrid Automaton (SHA) framework. Leveraging the fact that cell-biologic changes necessary for cancer development may be schematized as a series of discrete steps, an integrative closed-loop framework is proposed for describing the progressive development of cancer and determining optimal personalized therapies. The use of IPA techniques for optimal personalized cancer therapy design is introduced and a methodology applicable to stochastic models of cancer progression is developed. A case study of optimal therapy design for advanced prostate cancer is performed. The tradeoff between system optimality and robustness (or, equivalently, fragility) is explored so as to generate valuable insights on modeling and control of cancer progression.

The PhD dissertation completed by **Sepideh Pourazarm** is entitled "Control and Optimization Approaches for Energy-Limited Systems: Applications to Wireless Sensor Networks and Battery-Powered Vehicles". It studies control and optimization approaches to obtain energy-efficient and reliable routing schemes for battery-powered systems in network settings. First, incorporating a non-ideal battery model, the lifetime maximization problem for static wireless sensor networks is investigated. Adopting an optimal control approach, it is shown that there exists a time-invariant optimal routing vector in a fixed topology network. Furthermore, under very mild conditions, this optimal policy is robust with respect to the battery model used. Then, the lifetime maximization problem is investigated for networks with a mobile source node. Redefining the network lifetime, two versions of the problem are studied: when there exists no prior knowledge about the source node's motion dynamics and when the source node's trajectory is known in advance. For the former, the solution can be reduced to a sequence of nonlinear programming problems solved on line as the source node trajectory evolves. For the latter, an explicit off-line numerical solution is required. Second, the problem of routing vehicles with limited energy through a network with inhomogeneous charging nodes is studied. The goal is to minimize the total elapsed time, including traveling and recharging time, for vehicles to reach their destinations. Adopting a game-theoretic approach, the problem is investigated from two different points of view: user-centric vs. system-centric. The former is first formulated as a mixed integer nonlinear programming problem. Then, by exploiting properties of an optimal solution, it is reduced to a lower dimensionality problem. For the latter, grouping vehicles into subflows and including traffic congestion effects, a system-wide optimization problem is defined. Both problems are studied in a dynamic programming framework as well. Finally, the thesis quantifies the Price Of Anarchy (POA) in transportation networks using actual traffic data. The goal is to compare the network performance under user-optimal vs. system-optimal policies. First, user equilibria flows and origin-destination demands are estimated for the Eastern Massachusetts transportation network using speed and capacity datasets. Then, obtaining socially-optimal flows by solving a system-centric problem, the POA is estimated.

The PhD dissertation completed by **Xinmiao Sun** is entitled "Static and Dynamic Optimization Problems in Cooperative Multi-Agent Systems". It considers challenging static and dynamic problems encountered in cooperative multi-agent systems. First, a unified optimization framework is proposed for a wide range of tasks including consensus, optimal coverage, and resource allocation problems. It allows gradient-based algorithms to be applied to solve these problems. Gradient-based algorithms are shown to be distributed for a subclass of problems where objective functions can be decoupled. Second, the issue of global optimality is studied for optimal coverage problems where agents are deployed to maximize the joint detection probability. Objective functions in these problems are non-convex and no global optimum can be guaranteed by gradient-based algorithms developed to date. In order to obtain a solution close to the global optimum, the selection of initial condition is crucial. The initial state is determined by an additional optimization problem where the objective function is monotone submodular, a class of functions for which the greedy solution performance is guaranteed to be within a provable bound relative to the optimal performance. The bound is known to be within (1-1/*e*) of the optimal solution and is improved by exploiting the curvature information of the objective function. The greedy solution is subsequently used as an initial point of a gradient-based algorithm for the original optimal coverage problem. In addition, a novel method is proposed to escape a local optimum in a systematic way instead of randomly perturbing controllable variables away from a local optimum. Finally, optimal dynamic formation control problems are addressed for mobile leader-follower networks. Optimal formations are determined by maximizing a given objective function while continuously preserving communication connectivity in a time-varying environment. It is shown that in a convex mission space, the connectivity constraints can be satisfied by any feasible solution to a Mixed Integer Nonlinear Programming (MINLP) problem. For the class of optimal formation problems where the objective is to maximize coverage, the optimal formation is proven to be a tree which can be efficiently constructed without solving a MINLP problem. In a mission space constrained by obstacles, a minimum-effort reconfiguration approach is designed for obtaining the formation which still optimizes the objective function while avoiding the obstacles and ensuring connectivity.

5. INTERACTIONS/TRANSITIONS DURING REPORTING PERIOD

Participation/Presentations at Meetings, Conferences, Seminars

C.G. Cassandras gave invited talks/ plenary addresses/lectures at the following meetings/organizations:

- Inst. of Pure and Applied Mathematics, Los Angeles, CA, October 2015 (Invited Lecture).
- UC Berkeley, Institute of Transportation Studies, Berkeley, CA, October 2015 (Invited Seminar).
- Georgia Inst. of Technology, Dept. of Biology, Atlanta, GA, October 2015 (Invited Seminar)
- 5th IFAC Conference on Analysis and Design of Hybrid Systems, Atlanta, GA, October 2015 (Invited Talk).
- ORNL Workshop on Connected and Automated Vehicles: “The Road to the Future Urban Mobility”, Knoxville, TN, November 2015 (Invited Lecture).
- Boston University, Transportation Nudges: Experiments in Urban Mobility Conference, Boston, MA, November 2105 (Invited Panel Presentation).
- NSF Workshop on Smart Cities, Arlington, VA, December 2015 (Invited Lecture).
- Workshop on Smart Cities, Hiroshima, Japan, December 2015 (Invited Lecture).
- 54th IEEE Conference on Decision and Control, Osaka, Japan, December 2015

(Invited Talk).

- 54th IEEE Conference on Decision and Control, Osaka, Japan, December 2015 (Invited Panel Presentation).
- Ohio State University, Columbus, OH, April 2016 (Invited Seminar).
- Institute for Sustainable Energy, Boston, MA, April 2106 (Invited Panel Presentation)
- Chinese Academy of Science, Institute of System Science, Beijing, China, May 2016 (Invited Seminar).
- Symposium on Emerging Frontiers in Systems and Control, Beijing, China, May 2016 (Invited Lecture).
- Workshop on Control and Optimization of Discrete Event Systems, Nanjing, China, May 2016 (Invited Lecture).
- 13th Intl. Workshop on Discrete Event Systems, Xian, China, June 2016 (Invited Talk).
- National University of Singapore, Singapore, June 2016 (Invited Seminar).
- Nanyang Tech. University, Singapore, June 2016 (Invited Seminar).
- Workshop for High School Students and Teachers, 2016 American Control Conf., Boston, MA, July 2016 (Invited Lecture).
- 2016 American Control Conf., Boston, MA, July 2016 (Invited Talk).
- Michigan Institute of Data Science Symposium, Ann Arbor, MI, November 2016 (Invited Lecture).
- 55th IEEE Conf. on Decision and Control, Las Vegas, NV, December 2016 (Invited Talk).
- Workshop for High School Students and Teachers, 55th IEEE Conf. on Decision and Control, Las Vegas, NV, December 2016 (Invited Lecture).
- Workshop on Rich Data Backed Control and Optimization for Smart Cities, 55th IEEE Conf. on Decision and Control, Las Vegas, NV, December 2016 (Invited Lecture).
- 55th IEEE Conf. on Decision and Control, Las Vegas, NV, December 2016 (**Plenary Lecture**).
- Winter School on Intelligent Transportation, Nanyang Technological University, Singapore, January 2017 (Invited Lecture).
- Worcester Polytechnic Institute, Worcester, MA, March 2017 (Invited Seminar).
- Boston University, Boston, MA, April 2017 (Invited Lecture).
- Chinese Academy of Science, Beijing, China, May 2017 (Invited Seminar).
- Peking University, Beijing, China, May 2017 (Invited Seminar).
- Tsinghua University, Beijing, China, May 2017 (Invited Seminar).
- National University of Singapore, Singapore, June 2017 (Invited Seminar).
- Nanyang Technological University, Singapore, June 2017 (Invited Seminar).
- 7th oCPS School on Cyber-Physical Systems, Lucca, Italy, June 2017 (**Plenary Lecture**).
- 38th Intl. Petri Nets Conference, Zaragoza, Spain, June 2017 (**Plenary Lecture**).
- 20th IFAC World Congress, Toulouse, France (Invited Talk).
- 20th IFAC World Congress, Toulouse, France (**Plenary Lecture**).
- LAAS-CNRS, Toulouse, France, July 2017 (**Distinguished Lecture**).
- RISE program for High School Students, Boston University, Boston, MA, August 2017 (Invited Panel Presentation).
- Institute of Advanced Study Workshop, City University of Hong Kong, Hong Kong, August 2017 (Invited Lecture).
- Industry Event: Robotics and Automation in Smart Cities, Boston University, Boston, MA, October 2017 (Invited Lecture).
- Institute for Sustainable Energy, Boston, MA, October 2017 (Invited Lecture).
- MathWorks, Natick, MA, November 2017 (**Distinguished Lecture**).
- 56th IEEE Conference on Decision and Control, Melbourne, Australia, December 2017 (Invited Talk).

- Pacific Northwest National Laboratory, Richland, WA, January 2018 (**Distinguished Lecture**).
- Nanyang Technological University, Singapore, March 2018 (Three Invited Lectures).
- University of Cyprus KIOS Center, Nicosia, Cyprus, April 2018 (**Distinguished Lecture**).
- Peking University, Beijing, China, May 2018 (Invited Seminar).
- Chinese Academy of Science, Beijing, China, May 2018 (Invited Seminar).
- Huazhong University, Wuhan, China, May 2018 (Invited Seminar).
- Tsinghua University, May 2018 (Invited Seminar).
- DES Summer School, Sorrento, Italy, May 2018 (Invited lecture).
- 15th IFAC Symposium on Control in Transportation Systems, Savona, Italy, June 2018 (**Plenary Lecture**).
- IEEE, July 2018 (Invited Webinar).

Transitions
N/A

6. NEW DISCOVERIES, INVENTIONS, OR PATENT DISCLOSURES

Provisional patent application: “Method and System for Dynamic Parking Allocation in Urban Settings” Application Number 61/521,424; filed 8/9/2011 (still pending).

7. HONORS/AWARDS

C.G. Cassandras (Lifetime, selected):

- Lilly Fellow (1991), Kern Fellow (2012)
- Fellow of IEEE (1996), Fellow of IFAC (2008)
- 2011 IEEE Control Systems Technology Award
- IFAC Harold Chestnut Prize (1999)
- Distinguished Member Award, IEEE Control Systems Society (2006)
- Editor-in-Chief of *IEEE Transactions on Automatic Control* (1998-2009)
- President, IEEE Control Systems Society, 2012
- Zhu Kezhen Award, 2012
- 2014 Engineering Distinguished Scholar Award, Boston University

Honors/Awards received during grant period:

- Keynote/Plenary speaker at five international meetings/conferences
- Four Distinguished Lectures
- Elected Fellow of IETI (2018)