AFRL-AFOSR-VA-TR-2020-0123

Triggered Control for Distributed Optimization and Learning in Networked Multi-Agent Systems

Jorge Cortes
UNIVERSITY OF CALIFORNIA, SAN DIEGO
9500 GILMAN DR DEPT 621
LA JOLLA, CA 92093-0621

08/07/2020
Final Report

DISTRIBUTION A: Distribution approved for public release.

Air Force Research Laboratory
AF Office Of Scientific Research (AFOSR)/RTA2

DISTRIBUTION A: Distribution approved for public release.

Arlington, Virginia 22203
Air Force Materiel Command

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*
07/01/2020

2. REPORT TYPE
Final

3. DATES COVERED *(From - To)*
05/15/2015-05/14/2020

4. TITLE AND SUBTITLE
Triggered Control for Distributed Optimization and Learning in Networked Multi-Agent Systems

5a. CONTRACT NUMBER

5b. GRANT NUMBER
FA9550-15-1-0108

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
Dr. Jorge Cortes

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
University of California, San Diego
9500 Gilman Dr MC0934
La Jolla, CA 92093-0934

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
USAF, AFRL DUNS 143574726
AF Office of Scientific Research
875 North Randolph Street, RM 3112
Arlington, VA 22203-1954

10. SPONSOR/MONITOR'S ACRONYM(S)
AFOSR

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Distribution A. Approved for Public Release.

13. SUPPLEMENTARY NOTES

14. ABSTRACT
This proposal has laid down the foundation of a novel framework for the design of distributed triggered control strategies that endow networked systems with greater autonomy and decision making capabilities in dynamic environments subject to uncertainty and evolving task specifications. Our approach has combined the reactive nature of event-triggered control with the autonomous features of self-triggered control in providing algorithmic solutions to the scenarios of distributed optimization in cooperative networks and distributed learning under networked strategic interactions. The conceptual novelties of the project hinge upon the notion of agent abstractions and promises about future states, an original combination of event- and self-triggered information updates, new methods to reason and operate on set-valued information models, and new techniques for distributed controller design and analysis.

15. SUBJECT TERMS
Resource-aware control, distributed optimization, learning, network coordination, 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			Jorge Cortes
					19b. TELEPHONE NUMBER *(Include area code)*
				15	858-822-7930

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.

AFOSR Award FA9550-15-1-0108
Final report

"Triggered Control for Distributed Optimization and Learning in Networked Multi-Agent Systems"
PI: Jorge Cortés, UC San Diego

Project Objectives

This proposal seeks to lay down the foundation of a novel framework for the design of distributed triggered control strategies that endow networked systems with greater autonomy and decision making capabilities in dynamic environments subject to uncertainty and evolving task specifications. Our approach combines the reactive nature of event-triggered control with the autonomous features of self-triggered control in providing algorithmic solutions to the scenarios of distributed optimization in cooperative networks and distributed learning under networked strategic interactions. The focus on these scenarios is motivated by their broad applicability and the growing interest in optimizing the behavior of autonomous networks with limited resources while managing uncertainty about the state of the environment, the network, and the adversaries. The conceptual novelties of the project hinge upon the notion of agent abstractions and promises about future states, an original combination of event- and self-triggered information updates, new methods to reason and operate on set-valued information models, and new techniques for distributed controller design and analysis. The anticipated outcome is the synthesis of real-time coordination strategies that endow networked systems with superior decision making capabilities for coordinated action and greater robustness against the multiple sources of uncertainty and disruption present in real-world applications. The proposed framework will have a significant impact in current AFOSR missions by enabling the robust operation of assets in highly decentralized scenarios, where limited communication is available, interactions are evolving and subject to failures and adversarial attacks, and information is partial and dynamically evolving.

Research Results (2019-2020)

During the last year of the project, we made a number of significant contributions to the synthesis of resource-aware distributed controllers and the design of accelerated optimization algorithms for machine learning applications based on the idea of event-triggered control. We detail these contributions next.

Regarding **triggered approach to optimization and learning**, we have began to explore an exciting line of work that employs an opportunistic state-triggered approach to discretize accelerated continuous flows. The starting point of our approach is the recent body of exciting work seeks to shed light on the behavior of accelerated methods in optimization and machine learning via high-resolution differential equations. These differential equations are continuous counterparts of the discrete-time optimization algorithms, and their convergence properties can be characterized using the powerful tools provided by classical Lyapunov stability analysis. An outstanding question of pivotal importance is how to discretize these continuous flows while maintaining their convergence rates. We have pursued a novel approach to deal with this problem through the idea of opportunistic state-triggered control. Essentially, we take advantage of the Lyapunov functions employed to characterize the rate of convergence of high-resolution differential equations to design variable-stepsize forward-Euler discretizations that preserve the Lyapunov decay of the original dynamics. By design, these triggers ensure that the discretization retains the decay rate of the Lyapunov function. Since the evaluation of the Lyapunov function relies on knowledge of the problem optimizer, we rely on well-known bounds available for strongly-convex functions to synthesize triggers that do not require such knowledge. The flexibility of the proposed framework is not limited to forward-Euler discretizations and provides a promising path towards the understanding of the acceleration phenomenon and the design of new adaptive optimization algorithms.

Figure 1 shows a sample illustration of the performance of the proposed discretizations. *The paper (CP.8) contains a complete technical exposition of these contributions.*

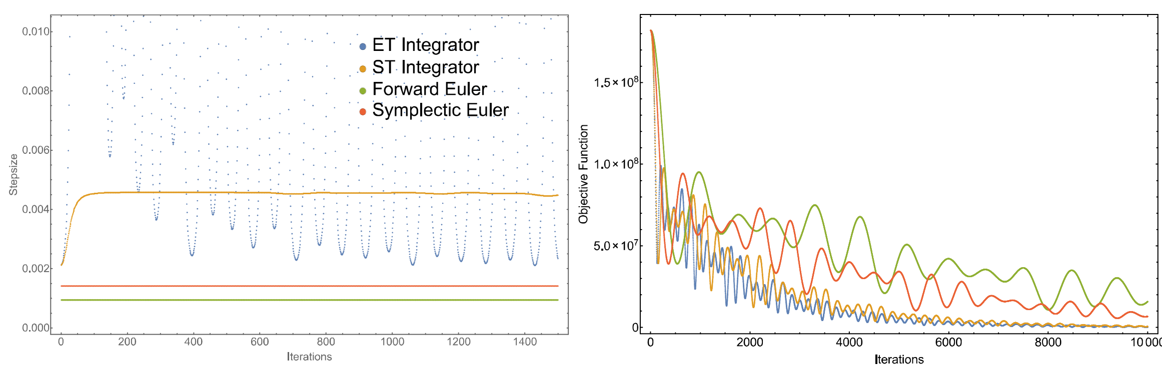

Figure 1: Triggered implementations of accelerated optimization flows. (a) compares the stepsizes along various discrete-time dynamics. The proposed discretizations (ET and ST integrators) keep a larger stepsize for the first 2000 iterations, approaching the optimizer much faster. (b) compares the evolution of the objective function along the different dynamics.

Regarding **triggered approach to controller design for safety-critical systems**, we have further developed with our study of previous years regarding the frequency behavior of nonlinear dynamics of power networks and, in particular, into guaranteeing the safe evolution of frequency transients in the presence of disturbances. Large frequency fluctuations can trigger generator relay-protection mechanisms and load shedding, which may further jeopardize network integrity, leading to cascading failures. Without appropriate operational architectures and control safeguards in place, the likelihood of such events is not negligible, given the high penetration of non-rotational renewable resources. These observations motivate us to develop control schemes to actively mitigate undesired transient frequency deviations under disturbances and contingencies. In previous reports, we have described our design of feedback controllers that meet the requirements of transient frequency safety and asymptotic stability. This controller was distributed and required no communication, in the sense that each control signal regulated on an individual bus only depends on neighboring system information that can be directly measured. However, the non-optimization-based nature of the proposed controller may cause bounded oscillations in the closed-loop system due to the lack of cooperation among control signals. To address this shortcoming, we have resorted to an MPC design implemented over a bilayer architecture. The bottom layer solves periodically a finite-horizon convex optimization problem and globally allocates control resources to minimize the overall control effort. The optimization problem incorporates a prediction model for the system dynamics, a stability constraint, and a relaxed frequency

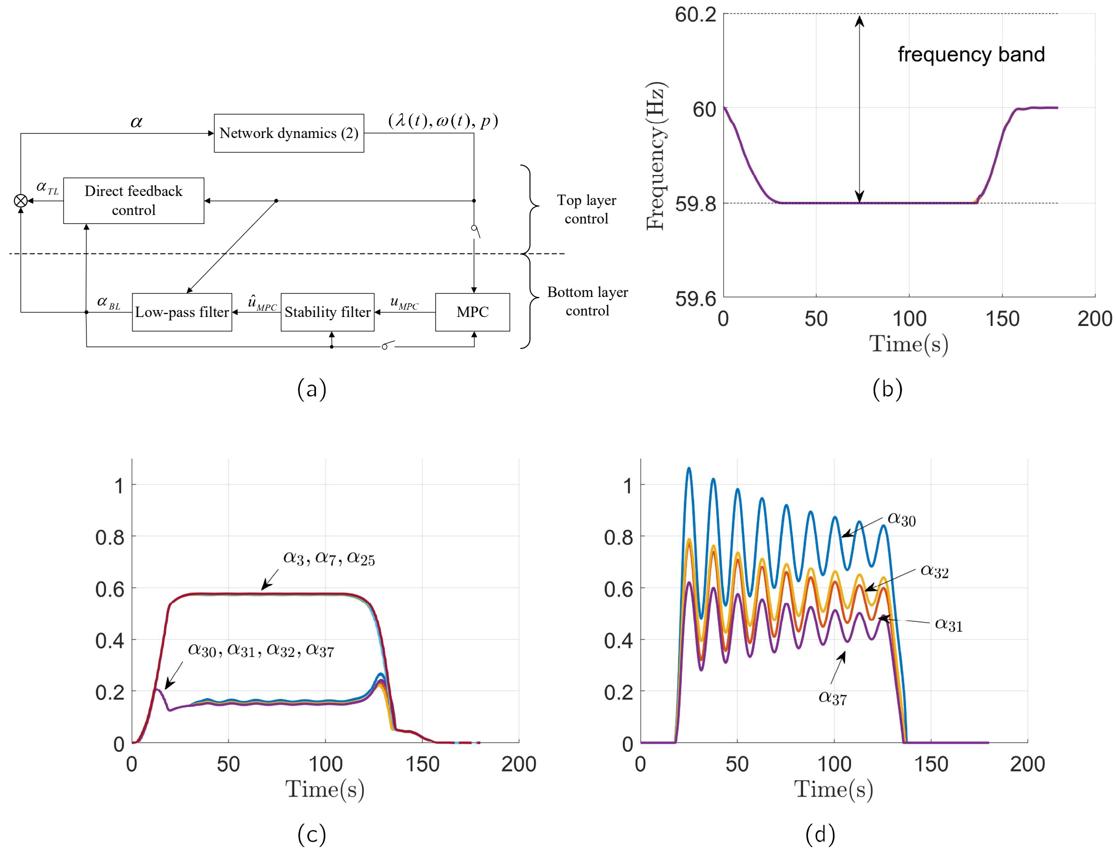

Figure 2: Performance of the proposed transient frequency controller on a IEEE-39 bus benchmark. (a) shows a block diagram of the closed-loop system with the proposed controller architecture; (b) shows the frequency responses at nodes 30, 31, 32, and 37 stay inside the safe region under the closed-loop system with the distributed controller; (c) shows the corresponding control input trajectories. Finally, (d) shows the oscillating control input trajectories obtained with our previous non-optimization-based strategy.

safety constraint. The prediction model is a linearized and discretized approximation of the nonlinear continuous-time power network dynamics, carefully chosen to preserve its local nature while keeping the complexity manageable. As a consequence, in the resulting convex optimization problem, the objective is the sum of local control costs, and each constraint only involves local decision variables. This enables us to apply saddle-point dynamics to recover its solution in a distributed fashion by allowing each bus (resp. line) to exchange system information within its neighboring buses (resp. lines). The top layer, as a real-time feedback controller, acts as a compensator, bridging the mismatch between the actual continuous-time power network dynamics and the sampled-based information employed in the bottom layer to rigorously guarantee frequency safety. The top layer control signal regulating on a generic bus only depends on physical measurements of system information within the range of its neighboring transmission lines.

Figure 4 shows a sample illustration of the performance of the proposed controller. *The paper (J.11) contains a complete technical exposition of these contributions.*

Research Results (2018-2019)

During this year, we continued our progress on the interplay between hybrid dynamics and nonlinear systems, resource-aware control and coordination, and the synthesis of distributed algorithmic solutions to optimization and learning problems. We next describe our progress.

Regarding **hybrid systems, triggered control, and distributed optimization**, we have studied the problem of integrating dynamic market policies with frequency regulation dynamics on similar timescales using the idea of resource-aware control and coordination. Power generation dispatch is typically done in a hierarchical fashion, where the different layers are separated according to their timescales. Broadly, at the top layer, economic efficiency is ensured via market clearing and at the bottom layer, frequency control and regulation is achieved via primary and secondary controllers. However, the intermittent and uncertain nature of distributed energy resources (DERs) and their integration represents a major challenge to the current design. Of particular concern is the need to maintain both frequency regulation and cost efficiency of regulation reserves in the face of increasing fluctuations in renewables. To this end, we have proposed an integrated dynamic market mechanism which combines the real-time market and frequency regulation, allowing competitive market players, including renewable generation, to negotiate electricity prices while using the most recent information on the frequency. A key aspect of our approach to achieve the integration of the different layers is the characterization of the robustness properties of the mechanisms used at each layer, since variables at the upper layers cannot be assumed in steady state any more at the lower ones. To analyze the stability of the resulting multi-rate hybrid interconnected system, we have identified a novel local Lyapunov function that includes the energy function of the closed-loop system. The availability of this function not only leads us to establish local exponential convergence to the desired equilibrium, but also allows us to rigorously establish its local input-to-state stability properties. Building on these results, we develop a time-triggered hybrid implementation that combines the discrete nature of iterative bidding with the continuous nature of the frequency evolution. In our design, we introduce two iteration loops, one (faster) inner-loop for the bidding process that incorporates at each step the frequency measurements, and one (slower) outer-loop for the market clearing and the updates in the power generation levels, that are sent to the continuous-time power network dynamics. We refer to this multi-rate hybrid implementation as time-triggered because we do not necessarily prescribe the time schedules to be periodic. To analyze its convergence properties, we regard the time-triggered implementation as an approximation of the continuous-time dynamics and invoke the robustness properties of the latter, interpreting their mismatch as a disturbance. This allows us to derive explicit upper bounds on the length between consecutive triggering times that guarantee that the time-triggered implementation remains asymptotically convergent. The computation of these upper bounds does not require knowledge of the efficient Nash equilibrium. Figure 3 presents a sample illustration of our results. *The paper (J.9) contains a complete technical exposition of these contributions.*

Regarding **triggered approach to controller design for safety-critical systems**, we have looked into the frequency behavior of nonlinear dynamics of power networks and, in particular, into guaranteeing the safe evolution of frequency transients in the presence of disturbances. During transients, power nodes are in danger of reaching their frequency limits and being tripped, which may in turn propagate throughout the interconnected system, causing blackouts. This phenomenon tends to happen more frequently in modern systems due to low inertia and highly-dynamic units. Therefore, there is a need to design controllers that ensure the safe evolution of the system. The technical approach to achieve

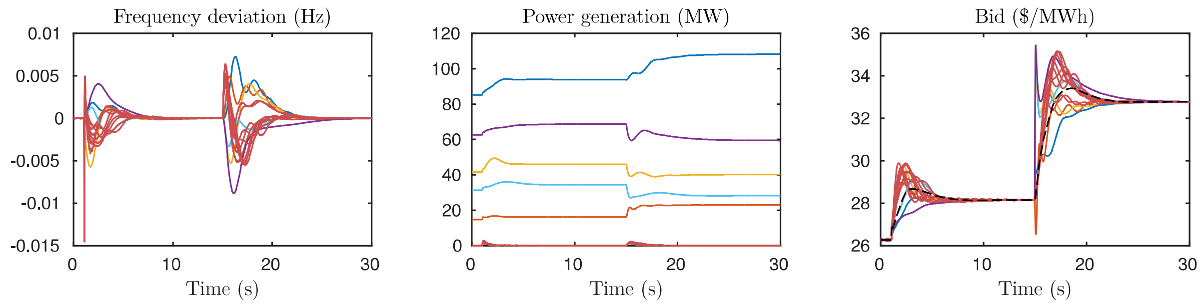

(a) Frequency deviations (b) Nodal power generations (c) Bids and Lagrange multiplier

Figure 3: Simulations of the interconnection on an IEEE-14 bus benchmark between the iterative bidding mechanism and power network dynamics modeled by the multi-rate hybrid system. After each sudden supply-demand imbalance, frequency is restored to its nominal value and the power generation quantities converge to the optimal values. Bids converge to the unique efficient Nash equilibrium.

this combines notions and tools from Lyapunov stability and set invariance theory, particularly the idea of encoding safe regions with barrier functions and employing these to design controllers that prescribe inputs gradually as the system evolution gets closer to violating safety. Our main result is the synthesis of a Lipschitz continuous, distributed controller, available at specific individual generator nodes, that renders the closed-loop system asymptotically stable; ensures that, for each controlled node, if its initial frequency belongs to a desired safe frequency region, its frequency trajectory stays in it for all subsequent time; and if, instead, its initial frequency does not belong to the safe region, then the frequency trajectory enters it in finite time, and once there, never leaves. For each one of these requirements, we provide equivalent mathematical formulations that are amenable to control design. Regarding stability, we consider an energy function and formalize it as identifying a controller that guarantees that the time evolution of this energy function along every trajectory of the dynamics is non-decreasing. Regarding invariance, we show that this condition is equivalent to having the controller make the safe frequency interval forward invariant. To avoid discontinuities in the controller design on the boundary of the invariant set, we resort to the notion of barrier functions to have the control effort gradually kick in as the state trajectory approaches the boundary. Our final step is to use the identified constraints to synthesize a specific controller that satisfies both and is distributed. The latter is a consequence of the fact that, for each bus, the constraints only involve the state of the bus and that of neighboring states. We analyze its robustness properties against measure error and parameter uncertainty, quantify its magnitude when the initial state is uncertain, and provide an estimation on the frequency convergence rate from the unsafe to the safe region for each controlled generator. Figure 4 shows a sample illustration of the performance of the proposed controller. *The papers (J.6) and (J.8) contain a complete technical exposition of these contributions.*

During this reporting period, we have also *published two tutorial papers (J.7) and (J.10)*, both of them building on the results obtained in this project. (J.7) covers event-triggered control and communication in network systems, with special emphasis on static consensus, and has been extremely well received for its expository and comprehensive nature. (J.10) is a tutorial paper on dynamic average consensus that seeks to fill a void in terms of providing coverage and a coherent exposition of a number of highly used tools and algorithms for control, estimation, and coordination that are scattered throughout the literature.

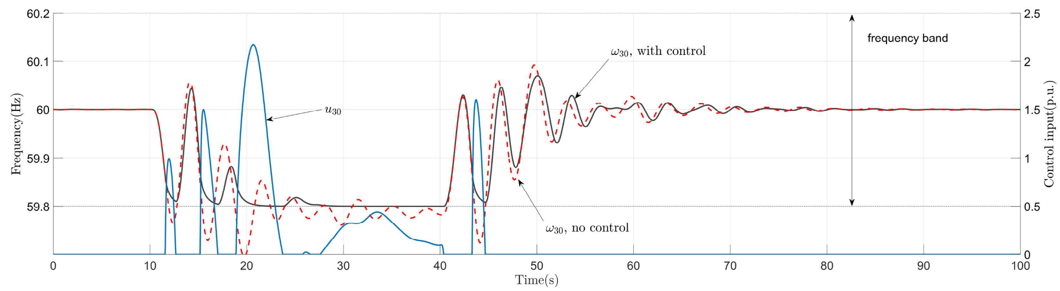

Figure 4: Performance of the proposed transient frequency controller on a IEEE-39 bus benchmark. Frequency and control input trajectories at node 30 corresponding to the power supply loss of a generator during [10,40]s. The frequency trajectory without transient controller goes beyond the safe bounds during the contingency, while this is avoided with the proposed controller. Note the latter only takes effect when the frequency is close to the safe bound.

Research Results (2017-2018)

During the third year of the project, we have continued with our progress on the synthesis of provably-correct distributed strategies, building on the opportunistic state-triggered control paradigm to improve efficiency, for large-scale optimization problems. We have also made significant progress on robust network optimization problems that incorporate uncertainty in their formulation and employ data to produce desirable out-of-sample performance guarantees. We next describe our progress along each of these themes.

Regarding **networked optimization and opportunistic state-triggered coordination**, we have obtained a major breakthrough on the characterization of the robustness properties of saddle-point dynamics, with significant implications for practical distributed implementations and online learning. Saddle-point dynamics and its variations are used extensively in the design and analysis of distributed feedback controllers and optimization algorithms in multiple domains, including networked cyberphysical systems, complex systems, and network flow problems. The study of the convergence properties of the saddle-point dynamics has traditionally focused on assessing the asymptotic stability of the equilibria. Unlike the case of gradient flows, the main roadblock is the absence of a natural certificate whose evolution is guaranteed to be strictly decreasing along the trajectories. This has motivated the use of either direct methods or LaSalle-type arguments. We have been able to identify a general form for strictly decreasing certificates for certain classes of constrained optimization problems. This result opens the way to the characterization of multiple additional properties beyond asymptotic stability, such as speed of convergence, disturbance rejection, and robustness to parameter and model uncertainty. In particular, we have established the input-to-state stability properties of the saddle-point dynamics, cf. Figure 5 for an illustration. This has also significant practical implications, including the robustness to communication link failures, the design of event-triggered strategies, and the development of performance guarantees for data-driven implementations that allow the network to incorporate measurements in a feedback, online fashion. *The paper (J.4) contains a complete technical exposition of these contributions.*

Regarding **robust optimization and distributed learning**, we have studied a class of stochastic optimization scenarios that finds numerous applications for multi-agent systems, such as target tracking, distributed estimation, and cooperative planning. Solving these problems, in an exact sense, requires

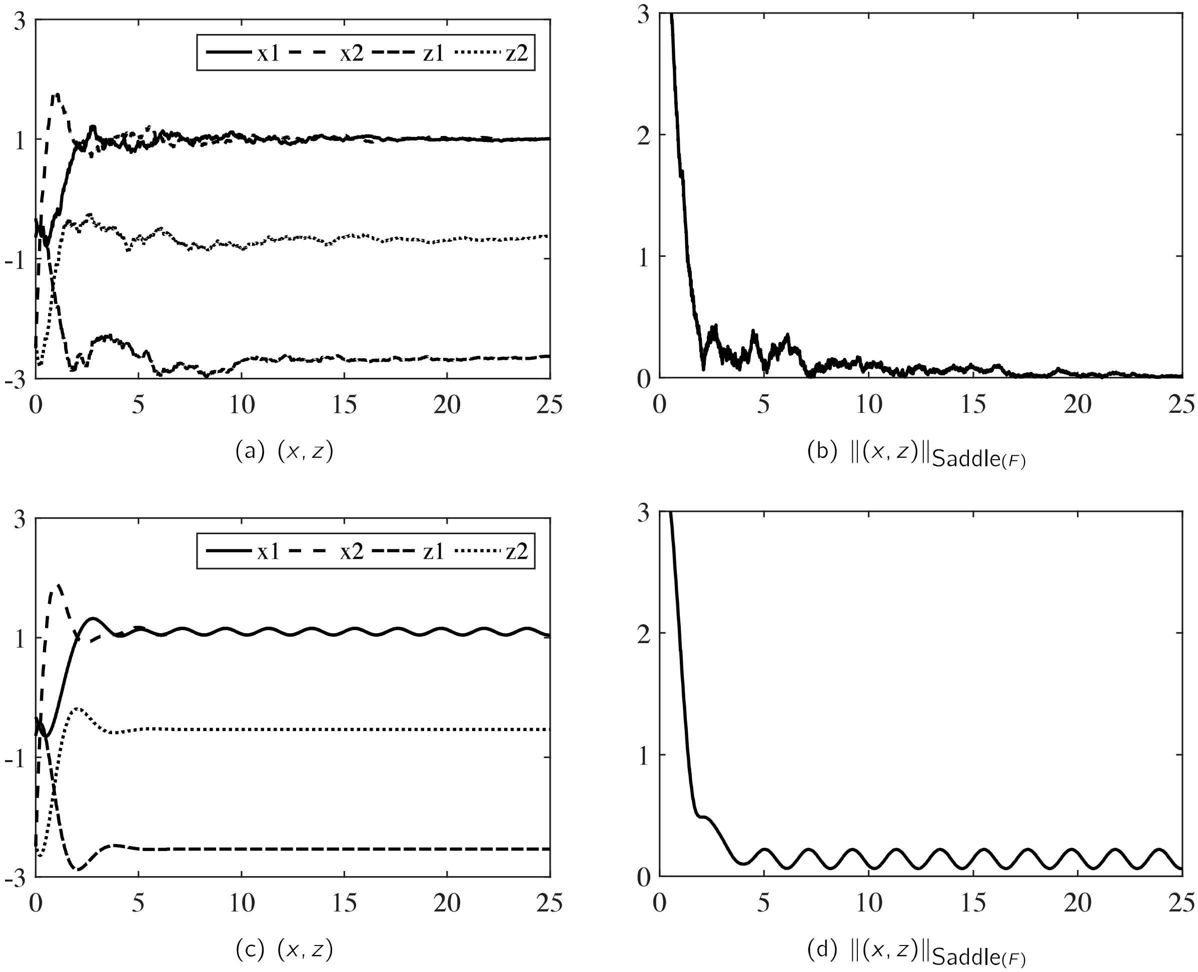

Figure 5: Illustration of the input-to-state stability properties of saddle-point dynamics for constrained, convex, network optimization problems. Performance is shown for both the case of an asymptotically vanishing disturbance, plots (a)-(b), and a persistent (constant plus a sinusoid) disturbance, plots (c)-(d).

the knowledge of the probability distribution of the random variables. Often this information is unavailable and instead, agents gather samples and use the data to find a solution to the stochastic optimization in an approximate sense. If the available dataset is large, machine learning algorithms are able to find the optimizer of the original problem with arbitrary precision. However, when the dataset is small, these algorithms fail to provide guarantees on the output obtained from the procedure. Scenarios with small datasets appear in applications where acquiring samples is expensive due to the size and complexity of the system or when decisions must be taken in real time, leaving less room for gathering many samples. Distributionally robust optimization (DRO), instead, uses finite datasets to provide a solution with desirable out-of-sample performance guarantees. Motivated by these observations, we have considered the task for a group of agents to collaboratively find a data-driven solution for a stochastic optimization problem using the tools provided by the DRO framework. Each agent could individually find a data-driven solution of the stochastic optimization. However, agents wish to cooperate to leverage on the data collected by everyone in the group. Instead of breaking the problem in two separate subproblems (model and learn first the uncertainty using data-fusion algorithms and then use the estimates to solve the stochastic optimization problem), the DRO method tackles them jointly, providing approximations of the optimizers tailored to the quality and size of the gathered data, making it amenable to online implementation.

We have reformulated the DRO problem over ambiguity sets defined as neighborhoods of the empirical distribution under the Wasserstein metric to display a structure amenable to distributed algorithm design. We have achieved this by augmenting the decision variables to yield a convex optimization whose objective function is the aggregate of individual objectives and whose constraints involve consensus among neighboring agents. Building on an augmented version of the associated Lagrangian function, we have identified a convex-concave function which under a min-max interchangeability condition has the property that its saddle-points are in one-to-one correspondence with the optimizers of the reformulated problem. Once in this form, we have designed and analyzed a distributed saddle-point dynamics for the identified convex-concave Lagrangian function. We have shown that the min-max interchangeability holds for two broad class of objective functions: the first class is the set of functions that are convex-concave in the decision and the random variable, respectively; the second class is where functions are convex-convex and have some additional structure: they are either quadratic in the random variable or they correspond to the loss function of the least-squares problem. *The paper (CP.4) contains a complete technical exposition of these contributions.*

Research Results (2016-2017)

During the second year of the project, we have continued with our progress on the synthesis of provably-correct distributed strategies for coupled optimization problems, extending our scope to a larger class of problems, and on the development of tools to characterize and quantify the robustness against disturbances in networked systems. We have also paid attention to the issue of protecting private data when executing coordination algorithms for networked optimization and have undertaken the design of opportunistic state-triggered controllers for nonlinear systems subject to actuator delays. We next describe our progress along each of these themes.

Regarding **networked optimization and distributed coordination**, we have considered general saddle-point problems with explicit agreement constraints on a subset of the arguments of both the convex and concave parts. These problems appear in dual decompositions of constrained optimization problems, and in other saddle-point problems where the convex-concave functions, unlike Lagrangians, are not necessarily linear in the arguments of the concave part. This is a substantial improvement over prior work that only focuses on dual decompositions of constrained optimization. For this type of problems, we have designed and analyzed distributed that are guaranteed to converge, scale well with the number of agents, and are robust against a variety of failures and uncertainties. The coordination algorithms that we study can be described as projected saddle-point subgradient methods with Laplacian averaging, which naturally lend themselves to distributed implementation. For these algorithms, we have characterized the asymptotic convergence properties in terms of the network topology and the problem data, and provided the convergence rate. *The paper (J.2) contains a complete technical exposition of these contributions.*

Regarding **privacy preservation** and **handling of data**, we have studied in depth the problem of solving networked optimization problems with privacy guarantees and high accuracy. We have settled on the notion of differential privacy because of its amenability to rigorous analytical treatment and its increasing technological impact. A system processing privacy sensitive inputs from individuals is made differentially private by randomizing its answers in such a way that the distribution over published outputs is not too sensitive to the data provided by any single participant. As a result, it is provably difficult for an adversary, no matter how powerful, to make inferences about individual records from the published outputs, or even to detect the presence of an individual in the dataset. Differential privacy is already being employed by Google and Apple in commercial products such as Chrome and the recently introduced HomePod. For the standard setup of a group of agents seeking to minimize the sum of their individual objective functions over a communication network, our first contribution shows that coordination algorithms which rely on perturbing the agents messages with noise cannot satisfy the requirements of differential privacy if the underlying noiseless dynamics are locally asymptotically stable. The presence of noise necessary to ensure differential privacy is known to affect the algorithm accuracy in solving the distributed convex optimization problem. However, this result explains why message-perturbing strategies incur additional inaccuracies that are present even if no noise is added. Figure 6 illustrates this point. Our second contribution is motivated by the goal

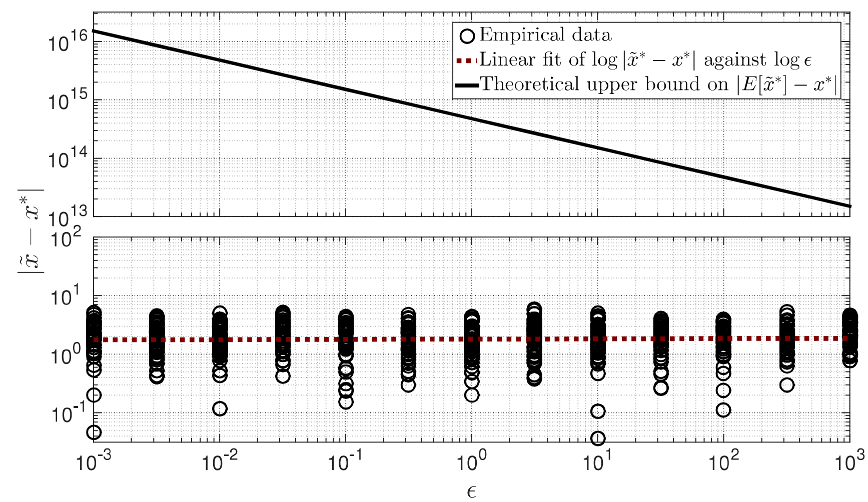

Figure 6: Privacy-accuracy trade-off for the algorithm proposed in "Z. Huang, S. Mitra, and N. Vaidya. Differentially private distributed optimization. In *Proceedings of the 2015 International Conference on Distributed Computing and Networking*, Pilani, India, January 2015" for $n = 10$ agents. With that paper's notation, the algorithm's parameters are $q = 0.1$, $p = 0.11$, $c = 0.5$. The stepsize $\alpha_k = cq^{k-1}$ has finite sum. The individual objective functions are two-dimensional quadratic functions defined over $D = [-5, 5]^2$. The circles, dotted line, and solid line illustrate simulation results for 50 executions, their best linear fit in logarithmic scale, and the theoretical upper bound on accuracy, respectively. The vertical axis is broken to better display the scale of the algorithm output.

of guaranteeing that the algorithm accuracy is only affected by the presence of noise. We propose a general framework for functional differential privacy over Hilbert spaces and introduce a novel definition of adjacency using adjacency spaces. The latter notion is quite flexible and includes, as a special case, the conventional bounded-difference notion of adjacency. We carefully specify these adjacency spaces within the L_2 -space such that the requirement of differential privacy can be satisfied with bounded perturbations. Our third contribution builds on these results on functional perturbation to design a class of distributed, differentially private coordination algorithms. We let each agent perturb its own objective function based on its desired level of privacy, and then the group uses any provably correct distributed coordination algorithm to optimize the sum of the individual perturbed functions. Two challenges arise to successfully apply this strategy: the fact that the perturbed functions might lose the smoothness and convexity properties of the original functions and the need to characterize the effect of the added noise on the minimizer of the resulting problem. We address the first challenge using a cascade of smoothening and projection steps that maintain the differential privacy of the functional perturbation step. We address the second challenge by explicitly bounding the absolute expected deviation from the original optimizer using a novel Lipschitz characterization of the argmin map. By construction, the resulting coordination algorithms satisfy the requirement of recovering perfect accuracy in the absence of noise. Figure 7 illustrates our design. *The paper (J.3) contains a complete technical exposition of these contributions.*

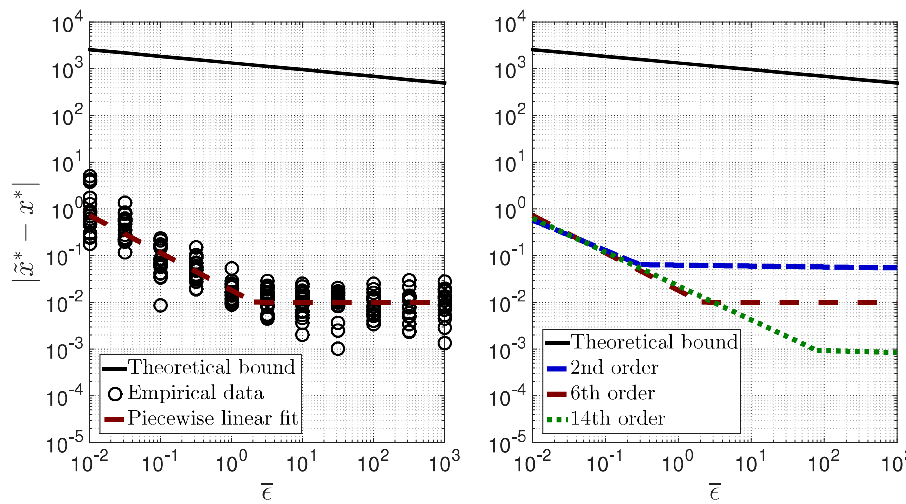

Figure 7: Privacy-accuracy trade-off curve of the proposed class of distributed, differentially private algorithms (with the same data as Figure 6) and different truncation orders. Left: empirical data and its best piecewise linear fit for 6th-order truncation of the function expansions, together with our theoretical upper bound. Right: piecewise linear fit of empirical data for 2nd, 6th, and 14th order truncations as well as the theoretical upper bound. As can be observed, accuracy can be improved arbitrarily by increasing the truncation order.

Regarding **opportunistic state-triggered coordination**, we have designed control strategies that stabilize nonlinear systems with time-varying delays in actuation. Our work has been motivated by the need for a more efficient use of the available resources and the consideration of increasingly realistic scenarios. The problem we have considered is particularly challenging due to the interplay between event-triggering and time delays: due to the opportunistic nature of event-triggering, the controller "waits" until the system tends to become unstable and then updates the control accordingly, but if this control takes some time to reach the system, it may no longer be able to prevent the system from instability. Therefore, the controller has to be "sufficiently more conservative" and update the control "sufficiently ahead in time" to ensure closed-loop stability, which makes the design challenging. For a wide class of nonlinear systems with arbitrarily large time delays, we have employed the method of predictor feedback to compensate for the system input delay to co-design the closed-loop control law and the triggering strategy to ensure the monotonic evolution of the Lyapunov function identified in the continuous-time case. For our triggering design, we have established global asymptotic stability of the closed-loop system and proved the existence of a uniform lower bound on the inter-event times, a result which rules out the possibility of Zeno behavior (i.e., the possibility of arbitrarily fast triggering leading to triggering times that accumulate at a finite time instant). We have further analyzed the particular case of a controllable linear system and give explicit expressions for the design variables. In the linear case, we also shown that the design achieves exponential stability. Our final contribution has been the analysis of the trade-off between communication cost and convergence speed in event-triggering control for the case of linear systems. *The paper (CP.2) contains a technical exposition of these contributions.*

Finally, regarding **robustness in networked systems**, we have studied the notion of transient stability in power networks. Transient stability refers to the ability of power networks to reach an acceptable state range while respecting operational constraints when subject to disturbances. Due to the uncertainty and variety of disturbances, it is challenging to characterize the disturbances under which power networks can still operate normally during transients. We have considered a linearized AC power network model with unknown amplitude and time-varying power disturbance injected at various buses. The generality and simplicity of this model makes our results applicable to networks that appear in other areas beyond power systems. We term a disturbance feasible for transient stability if, under such disturbance and for a desired period of time, the frequency of each bus and the power flow in each transmission line still remain in their respective bounds. Our main goal has been the design of efficient ways of computing the transient-state feasibility set consisting of all such disturbances. Since this set contains infinitely many constraints, we have developed a sampling method to approximate it by identifying a subset and a superset. We compute the subset approximation by sampling and tightening the constraints at finite discrete-time instants, and using the linear dynamics of the network to upper and lower bound the evolution of state signals, so as to ensure that all constraints are respected at all times. The superset approximation comes from using only a finite number of the constraints appearing in the exact set. We show that the approximation sets converge to the real transient-state feasibility set as we increase the number of sampling points. We have also defined a metric to measure the approximation conservativeness by estimating the region difference between the approximations and the real set. Finally, we have designed an algorithm to optimize this metric by adjusting the positions of the sampling points in a way that monotonically decreases the value of the error metric. *The paper (CP.3) contains a technical exposition of these contributions.*

Research Results (2015-2016)

During the first year of the project, our specific objectives were to make inroads in the synthesis of provably-correct distributed strategies to solve coupled optimization problems, to explore to what extent such strategies could be implemented in real time using ideas from opportunistic state-triggered control, and finally to characterize the robustness properties that such algorithms provide to the networked multi-agent system. For these objectives, we have obtained significant results that we describe next.

Regarding **distributed optimization** and **opportunistic state-triggered coordination**, we have designed provably correct distributed dynamics which, together with a set of distributed criteria to trigger state broadcasts among neighbors, enable a group of agents to collectively solve linear programs in standard form. The algorithm performance is illustrated in a multi-agent assignment problem in Figure 8. Our starting point has been the introduction of a novel distributed continuous-time dynamics for

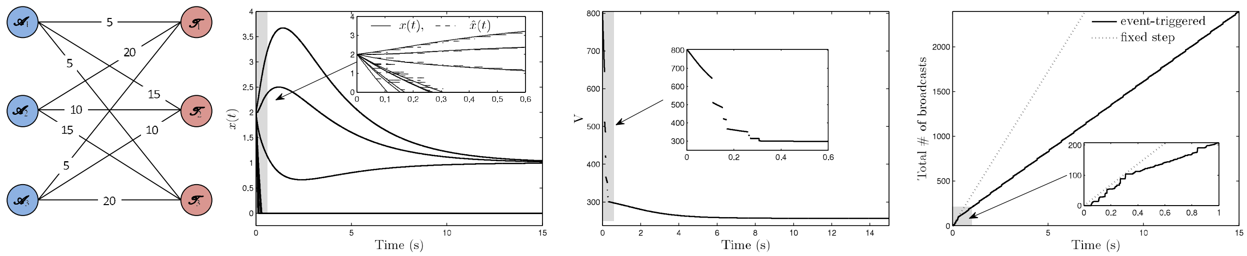

(a) Agent-task graph (b) Broker state trajectories (c) Lyapunov function (d) Cumulative broadcasts

Figure 8: Illustratation of the proposed event-triggered coordination algorithm in a multi-agent assignment. This is a resource allocation problem where N tasks are to be assigned to N agents. Each potential assignment of a task to an agent has an associated benefit and the global network objective is to maximize the sum of the benefits in the final assignment. The assignment of an agent to a task is managed by a broker. (a) shows the assignment graph with agents $\mathscr{A}_i$ in blue, tasks $\mathscr{T}_j$ in red, and the benefit of a potential assignment as edge weights. (b) shows the state trajectories of the brokers. The brokers' state is $x = (x_{1,1}, x_{1,2}, x_{1,3}, x_{2,1}, x_{2,2}, x_{2,3}, x_{3,1}, x_{3,2}, x_{3,3})$ and the inlay also shows the evolution of the broadcast states, $\hat{x} = (\hat{x}_{1,1}, \hat{x}_{1,2}, \hat{x}_{1,3}, \hat{x}_{2,1}, \hat{x}_{2,2}, \hat{x}_{2,3}, \hat{x}_{3,1}, \hat{x}_{3,2}, \hat{x}_{3,3})$ in dashed lines. The aggregate of the brokers' states converge to the unique solution $\{(0, 1, 0, 1, 0, 0, 0, 0, 1)\}$. The Lyapunov function V is discontinuous but decreasing, as evidenced in (c). The cumulative number of broadcasts appears roughly linear and the execution is persistently flowing. The dashed line in (d) illustrates the cumulative number of broadcasts required for a fixed-stepsize implementation of the dynamics.

linear programming based on an exact quadratic regularization and the characterization of its solutions as saddle points of an augmented Lagrangian function. This distributed dynamics is discontinuous in the agents' state because of the inequality constraints in the original linear program. Our approach to synthesize strategies that rely only on discrete-time communication then proceeds by having agents implement the distributed continuous-time dynamics using a sample-and-hold value of their neighbors state.

The key challenge is then to identify suitable criteria to opportunistically determine when agents should share information with their neighbors in order to guarantee asymptotic convergence and persistency of the executions. Because of the technical complexity involved in solving this challenge, we divide our technical developments in two steps, dealing first with the design of *centralized criteria* and then *distributed criteria*. Under our centralized event-triggered communication scheme, agents use global knowledge of the network to determine when to synchronously broadcast their state. The characterization of the convergence properties of the centralized implementation is challenging because the original continuous-time dynamics is discontinuous in the agents state and the fact that its final convergence value (being the solution of the optimization problem) is not known a priori, which further complicates the identification of a common smooth Lyapunov function. Using concepts and tools from switched and hybrid systems, we overcome these obstacles by introducing a discontinuous Lyapunov function and examining its evolution during time intervals where state broadcasts do not occur. We build on our centralized design to synthesize a distributed event-triggered communication law under which agents use local information to determine when to broadcast their individual state. Our strategy to accomplish this is to investigate to what extent the centralized triggers can be implemented in a distributed way and modify them when necessary. In doing so, we face the additional difficulty posed by the fact that the mode switches associated to the discontinuity of the original dynamics are not locally detectable by individual agents. To address this challenge, we bound the evolution of the Lyapunov function under mode mismatch and, based on this understanding, design the distributed triggers so that any potential increase of the Lyapunov function due to the mismatch is compensated by the decrease in its value before the mismatch took place. Moreover, the distributed character of the agent triggers leads to asynchronous state broadcasts, which poses an additional challenge for both design and analysis. Our main result establishes the asymptotic convergence of the distributed implementation and identifies sufficient conditions for executions to be persistently flowing (that is, state broadcasts are separated by a uniform time infinitely often). We show that the asynchronous state broadcasts cannot be the cause of non-persistently flowing executions and we conjecture that all executions are in fact persistently flowing.

Regarding the **robustness characterization** of networked multi-agent systems operating under the proposed algorithms, we have shown, as a byproduct of using a hybrid systems modeling framework in our technical approach, that the global asymptotic stability of the proposed distributed coordination strategy for linear programming is robust to small perturbations. Motivated by the need to develop metrics and tools to measure the robustness of multi-agent systems, in particular power networks, we have worked on studying their vulnerability under voltage and frequency fluctuation. Despite the careful design of power networks, large blackouts still happen due to various factors, including relatively small power flow deviations after disconnecting some transmission lines and loss of frequency synchronization leading to cascading failures. We have defined a notion of robustness metric for a power network in terms of the minimal power disturbance required to cause an initial failure. In our approach, we allow for multiple disturbances to act upon different load-side nodes. We have formulated the computation of the robustness metric in terms of various equivalent optimization problems with affine inequality constraints. To obtain these results, our analysis has identified a polyhedral cone determined by the topology of the power network whose extreme rays determine the affine inequality constraints.

Finally, we have designed a novel **event-triggered broadcasting and controller** update strategy to solve the multi-agent average consensus problem over weight-balanced digraphs. Due to its numerous applications in networked systems, many algorithmic solutions exist to the multi-agent average consensus problem; however, a majority of them rely on agents having continuous or periodic availability of information from other agents. Unfortunately, this assumption leads to inefficient implementations in terms of energy consumption, communication bandwidth, congestion, and processor usage. Our proposed law does not require individual agents to have continuous access to information about the state of their neighbors and is fully distributed in the sense that it does not require any a priori knowledge by agents of global network parameters to execute the algorithm. Our Lyapunov-based design builds on the evolution of the network disagreement to synthesize triggers that agents can evaluate using locally available information to make decisions about when to broadcast their current state to neighbors. In our design, we carefully take into account the discontinuities in the information available to the agents caused by broadcasts received from neighbors and their effect on the feasibility of the resulting implementation. Our analysis shows that the resulting asynchronous network executions are free from Zeno behavior, i.e., only a finite number of events are triggered in any finite time period, and exponentially converge to agreement on the average of all agents initial states over weight-balanced, strongly connected digraphs. We also provide a lower bound on the exponential convergence rate and characterize the asymptotic convergence of the network under switching topologies that remain weight-balanced and are jointly strongly connected. Lastly, we have proposed a periodic implementation of our event-triggered design that has agents check the triggers periodically and characterize the sampling period that guarantees correctness.

Research Outcomes

Refereed Journal Articles

(J.1) D. Richert and J. Cortés. Distributed linear programming with event-triggered communication. *SIAM Journal on Control and Optimization*, 54(3):1769–1797, 2016

(J.2) D. Mateos-Núñez and J. Cortés. Distributed saddle-point subgradient algorithms with Laplacian averaging. *IEEE Transactions on Automatic Control*, 62(6):2720–2735, 2017

(J.3) E. Nozari, P. Tallapragada, and J. Cortés. Differentially private distributed convex optimization via functional perturbation. *IEEE Transactions on Control of Network Systems*, 5(1):395–408, 2018

(J.4) A. Cherukuri, E. Mallada, S. H. Low, and J. Cortés. The role of convexity in saddle-point dynamics: Lyapunov function and robustness. *IEEE Transactions on Automatic Control*, 63(8):2449–2464, 2018

(J.5) J. Cortés and M. Egerstedt. Coordinated control of multi-robot systems: A survey. *SICE Journal of Control, Measurement, and System Integration*, 10(6):495–503, 2017

(J.6) Y. Zhang and J. Cortés. Distributed transient frequency control for power networks with stability and performance guarantees. *Automatica*, 105:274–285, 2019

(J.7) C. Nowzari, E. Garcia, and J. Cortés. Event-triggered control and communication of networked systems for multi-agent consensus. *Automatica*, 105:1–27, 2019

(J.8) Y. Zhang and J. Cortés. Characterizing tolerable disturbances for transient-state safety in power networks. *IEEE Transactions on Network Science and Engineering*, 6:210–224, 2019

(J.9) T. Stegink, A. Cherukuri, C. De Persis, A. J. van der Schaft, and J. Cortés. Hybrid interconnection of iterative bidding and power network dynamics for frequency regulation and optimal dispatch. *IEEE Transactions on Control of Network Systems*, 6(2):572–585, 2019

(J.10) S. S. Kia, B. Van Scoy, J. Cortés, R. A. Freeman, K. M. Lynch, and S. Martinez. Tutorial on dynamic average consensus: The problem, its applications, and the algorithms. *IEEE Control Systems*, 39(3):40–72, 2019

(J.11) Y. Zhang and J. Cortés. Distributed bilayered control for transient frequency safety and system stability in power grids. *IEEE Transactions on Control of Network Systems*, 2020. To appear

Book Chapters

(BC.1) C. Nowzari, J. Cortés, and G. J. Pappas. Event-triggered communication and control for multi-agent average consensus. In Y. Wang, E. Garcia, D. Casbeer, and F. Zhang, editors, *Cooperative Control of Multi-Agent Systems: Theory and Applications*, pages 177–207. Wiley, New York, 2017

Refereed Conference Proceedings

(CP.1) Y. Zhang and J. Cortés. Quantifying the robustness of power networks against initial failure. In *European Control Conference*, pages 2072–2077, Aalborg, Denmark, July 2016

(CP.2) E. Nozari, P. Tallapragada, and J. Cortés. Event-triggered control for nonlinear systems with time-varying input delay. In *IEEE Conf. on Decision and Control*, pages 495–500, Las Vegas, NV, 2016

(CP.3) Y. Zhang and J. Cortés. Transient-state feasibility set approximation of power networks against disturbances of unknown amplitude. In *American Control Conference*, pages 2767–2772, Seattle, WA, May 2017

(CP.4) A. Cherukuri and J. Cortés. Data-driven distributed optimization using Wasserstein ambiguity sets. In *Allerton Conf. on Communications, Control and Computing*, pages 38–44, Monticello, IL, 2017

(CP.5) Y. Zhang and J. Cortés. Distributed transient frequency control in power networks. In *IEEE Conf. on Decision and Control*, pages 4595–4600, Miami Beach, FL, December 2018

(CP.6) Y. Zhang and J. Cortés. Transient frequency control with regional cooperation for power networks. In *IEEE Conf. on Decision and Control*, pages 2587–2592, Miami Beach, FL, December 2018

(CP.7) Y. Zhang and J. Cortés. Double-layered distributed transient frequency control with regional coordination. In *American Control Conference*, pages 658–663, Philadelphia, PA, July 2019

(CP.8) M. Vaquero and J. Cortés. Convergence-rate-matching discretization of accelerated optimization flows through opportunistic state-triggered control. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32, pages 9767–9776. Curran Associates, Inc., 2019