Summary of Research Contributions

Mahnoosh Alizadeh

My research is broadly focused on mathematical modeling and analysis of multi-agent **Human-Cyber-Physical** Systems (H-CPS). I am particularly interested in designing scalable data-driven control and optimization frameworks and market mechanisms for societal-scale infrastructure systems. This includes smart power grids, electric transportation systems, and energy-aware communities and buildings. H-CPS are networked systems of physical assets with embedded intelligence and human-machine interfaces. Ubiquitous intelligence and information exchange and the active involvement of humans in the control loop distinguish them from their predecessors and gives them the potential to be more adaptive, resilient, and safe. However, these new characteristics also make our future infrastructure vulnerable because no single engineering or science discipline has all the right tools to design simple and scalable monitoring and control solutions to manage their operations reliably and efficiently. My research interests lie in addressing this broad interdisciplinary challenge, where I often draw and build on theory from several disciplines including stochastic optimization, learning theory, controls, power systems engineering, transportation engineering, statistical signal processing, and game theory.

Next, **I summarize selected prior contributions** under four categories, which are presented in chronological order starting from my initial contributions in my PhD thesis.

Large-scale monitoring and control of electricity demand

Enabling demand response (DR) technologies to harness the intrinsic flexibility of end-use electricity demand (load) is an inevitable step for efficient trading of high levels of non-dispatchable clean generation resources. For the residential and commercial sectors, DR poses a highly complex problem that suffers from a curse of dimensionality. Electricity load is comprised of a large number of heterogeneous subcomponents which represent the load of appliances that are gradually plugged into the grid by customers with heterogeneous preferences. The main focus of my PhD thesis is on reduced-order modeling frameworks for scalable control and monitoring of large populations of electric appliances and accompanying incentive mechanisms with differentiated service tiers for customer participation in direct control programs (see, e.g., [1,2]). The approach is unique in that it does not discard the individual constraints of different appliances when approximating aggregate load flexibility. Instead, we used a basic idea that is common in scalable multimedia applications for *graceful degradation*: similar electricity consumption requests can be effectively bundled in clusters and their aggregate flexibility can be captured by a representative linear state-space model. This approach enables highly reduced computational and communication requirements for load modeling, market planning and real-time control with adjustable margins of error (a tradeoff between scalability and reliability). Moreover, differentiated (a.k.a., nonlinear) pricing schemes allow customers with higher degrees of flexibility to pay less for their electric energy consumption. Lastly, our approach allows for anonymous interaction between customers and the electricity retailer, a factor that can play a vital role in gaining customer trust for a widespread implementation of direct load scheduling programs. Our method covers various load classes such as thermostatically controlled loads [3,4] and electric vehicles [5]. We have also investigated the benefits of such reduced order models for market integration of DR [6,7].

Electric vehicle management in coupled power and transportation systems

Our societal infrastructures are complex interdependent networked systems. For example, since their inception, power systems have depended on the delivery of fuels to generation plants through transportation networks. Similar connections can be found between power networks and gas networks, data networks, water networks, emergency response systems, etc. Such interdependencies are getting closer to real-time operations as we increase stochasticity on the supply side and rely on DR to increase the use of green energy in our power grids. Specifically, I am interested in the fact that by introducing new control measures such as DR mechanisms to shift energy use in our societal infrastructure to times and locations at which energy is cheap, we introduce real-time feedback loops between various processes and services that are potentially handled by separate organizations and have been optimized independently until now. Hence, we are further increasing their control system design complexity both in terms of analytical modeling and optimization as well as economic mechanism design to align the interests of the many organizations involved.

Power and transportation networks are arguably the quintessential coupled infrastructure system in which all of the above challenges are present. I am especially interested in the role of transportation electrification in coupling power and transportation system operations. Designing a reliable Electric Vehicle (EV) smart battery charging scheme requires considering the fact that EVs are primarily modes of transportation. In the past six years, I have dedicated significant effort to design and implementation of **mobility-aware routing and charging mechanisms for electric vehicles**. The problem combines features of large-scale optimization, optimal control, traffic routing, game theory, multi-agent system theory, and power system engineering - hence, modeling complexity is at its heart. The paper [8], which was published in TCNS, was the first to mathematically model the spatial coupling effects between power and transportation networks due to electric vehicle mobility and propose distributed coordination and pricing mechanisms that can drive both networks towards a socially optimal operational point. A subsequent publication in TCNS [9] explores the temporal dynamics of the coupled power and transportation networks problem in the context of charging service providers who wish to eliminate queues at their charging stations. Lastly, a third publication in TCNS [10] proposes computational methods for the optimal coordination of autonomous mobility-on-demand (AMoD) systems, with a central focus on accounting for the couplings between the power and transportation networks. AMoD is a rapidly developing mode of transportation wherein self-driving, electric vehicles transport passengers on demand (similar to an Uber system, but with self-driving, electric vehicles). The ability to intelligently route empty vehicles will provide a unique opportunity for joint traffic management and energy dispatch. For example, autonomous electric vehicles could act as distributed mobile storage devices, providing a support mechanism for integration of intermittent renewable energy resources and relieving congestion in energy distribution network. While [10] adopts a welfare-maximizing model for the AMoD service provider, our most recent TCNS paper [11] studies the effects of having profit-maximizing AMoD service providers under monopolistic and duopolistic settings on customer welfare and ride prices. Note that since the above-mentioned papers all adopt a deterministic network flow formulation, they cannot generate real-time solutions or handle stochasticities. As such, we have further studied real-time algorithms that consider stochasticity of network conditions, integrality constraints in vehicle routing decisions, or lack of accurate statistics about system parameters, e.g., [12–14]. Akin to my work on differentiated pricing mechanisms for direct load control, we have also proposed differentiated routing and pricing algorithms that enable direct allocation of users to electric vehicle fast charging stations in order to efficiently operate a network of stations for welfare or profit maximization [15]. Moreover, in addition to my theoretical studies on this problem, I have also been actively involved in real-world implementations of smart charging algorithms at workplace charging facilities at the Google Mountain View campus, SLAC national lab, and Stanford University's shuttle system.

Bandit learning for cyber-physical systems

Stochastic bandit optimization algorithms have long found applications in fields where some characteristics of the users' response are not known and can only be learned through a limited number of queries with noisy observations, while controlling cumulative reward gained during the learning process. This is a very common challenge in H-CPS due to the involvement of humans in the loop. For example, bandit optimization can naturally be employed towards increasing efficiency in societal systems where some characteristics of user response are not known and has to be learned in a few interactions, e.g., to optimally price electricity or dispatch a population of demand response assets in power systems. However, one cannot directly adopt existing bandit optimization algorithms for such safety-critical networks. For example, in the case of data-driven optimization of demand response or electricity pricing, dispatch signals posted by the bandit algorithm might violate distribution system constraints as the response of users to dispatch signals are not fully known and cannot be deterministically accounted for in the distribution system power flow constraints. To address this challenge, my goal is to derive **regret-optimal bandit optimization policies that would honor a system's safety and reliability requirements** in spite of uncertainty about system parameters. Several publications including papers published in NeurIPS 2019 and 2020 [16–18] study the effect of **safety constraints in linear stochastic bandits** and provide theoretical regret guarantees for several new safe learning algorithms (extensions to Gaussian Processes considered in [19]). While such linear settings are admittedly much simpler than those required by most H-CPS applications, they provide intuition on the basic principles of devising safe algorithms that can efficiently balance exploration and exploitation. Generalizing these insights, our work in [20], published in TSG, has also explored the application of **safe bandit optimization for managing electricity demand** with unknown price response from users, while respecting distribution system constraints. More recently, we have been investigating the effect of multiple decision makers common to many CPS in the design of multi-agent bandit algorithms [21], as well as model selection methods that start with multiple hypotheses about user parameters and adaptively learn the best model for linear bandit problems [22]. In particular, the agent has (potentially erroneous) information about the models but does not know the identity of the one(s) that the linear bandit task has been selected from. The goal of the agent is to identify the true model and transfer its collected experience to speedup the learning of the task at hand.

Distributed control and decision making in adversarial environments

Multi-agent systems commonly rely on the collective behavior of decision-makers distributed over a network, which can make them vulnerable to various forms of adversarial influences. I am interested in characterizing equilibria in such adversarial environments and designing **algorithmically-robust** learning and optimization mechanisms to protect against such vulnerabilities. Specifically, focused on distributed resource allocation through **network utility maximization** (NUM), our first line of work considers how man-in-the-middle Byzantine attacks can lead to non-convergence and infeasibility issues. Our work in TCNS [23] proposes attack-resilient primal-dual algorithms which rely on robust mean estimation and gradient averaging techniques to address these issues. In a related line of work, [24–26] consider the effect of probabilistically corrupt gradients on the convergence of gradient descent algorithms in the potentially broader scope of distributed optimization algorithms.

Considering a similar theme of distributed adversarial influence in the context of **graphical coordination games** (a game where agents in a network choose between two conventions and derive benefits from coordinating neighbors), our TCNS paper [27] studies the effects that system-level information and complexity of strategy space have on an adversary's ability to degrade aggregate welfare. Furthermore, to defend against such attacks in graphical coordination games, our publication in TAC [28] studies how a system operator can tweak log-linear learning algorithms to defend against 1) broad attacks, where the adversary distributes targeted incentives to all agents in the network to promote incoordination and 2) focused attacks, where the adversary can force a selected subset of the agents to commit to a prescribed convention without regards to their neighbors' actions. The paper's main contribution is to characterize the operator's fundamental trade-off between security against worst case broad attacks and vulnerability from focused attacks.

Last but not least, we have recently considered the challenges of distributed decision making in such adversarial environments in a model of competitive resource allocation known as the **Colonel Blotto game** (or its close variant, General Lotto), where two or more Colonels compete over control of multiple battlefields. Our work in [29] characterizes the impact of the division of resources among multiple subplayers competing against the same adversary on the best-case efficiency of the resulting collective behavior, where the main result unintuitively shows that a more balanced division of resources between subplayers can offer better performance guarantees than a more centralized allocation. Focused on quantifying the value of information in competitive resource allocation, papers [30, 31] consider the effect of asymmetric uncertainties about the players' resource endowments or battlefield valuations and how they could offer strategic advantages to one player.

References

[1] M. Alizadeh, A. Scaglione, A. Applebaum, G. Kesidis, and K. Levitt, "Reduced-order load models for large populations of flexible appliances," *IEEE Transactions on Power Systems*, vol. 30, no. 4, pp. 1758–1774, 2014.

[2] M. Alizadeh, Y. Xiao, A. Scaglione, and M. Van Der Schaar, "Dynamic incentive design for participation in direct load scheduling programs," *IEEE Journal of Selected Topics in Signal Processing: Special Issue on Smart Power Grids*, vol. 8, no. 6, pp. 1111–1126, 2014.

[3] M. Alizadeh and A. Scaglione, "Least laxity first scheduling of thermostatically controlled loads for regulation services," in *IEEE Global Conference on Signal and Information Processing*, 2013, pp. 503–506.

[4] K. Hreinsson, A. Scaglione, and M. Alizadeh, “An aggregate model of the flexible energy demand of thermostatically controlled loads with explicit outdoor temperature dependency,” in *HICSS-53*, 2020.

[5] M. Alizadeh, A. Scaglione, J. Davies, and K. S. Kurani, “A scalable stochastic model for the electricity demand of electric and plug-in hybrid vehicles,” *IEEE Transactions on Smart Grid*, vol. 5, no. 2, pp. 848–860, 2013.

[6] C. Chen, S. Kishore, Z. Wang, M. Alizadeh, and A. Scaglione, “A Cournot game analysis on market effects of queuing energy request as demand response,” in *IEEE Power and Energy Society General Meeting*, 2012.

[7] K. Hreinsson, A. Scaglione, M. Alizadeh, and Y. Chen, “New insights from the Shapley-Folkman lemma on dispatchable demand in energy markets,” *IEEE Transactions on Power Systems*, to appear, 2021.

[8] M. Alizadeh, H.-T. Wai, M. Chowdhury, A. Goldsmith, A. Scaglione, and T. Javidi, “Optimal pricing to manage electric vehicles in coupled power and transportation networks,” *IEEE Transactions on Control of Network Systems*, vol. 4, no. 4, pp. 863–875, 2016.

[9] M. Alizadeh, H.-T. Wai, A. Goldsmith, and A. Scaglione, “Retail and wholesale electricity pricing considering electric vehicle mobility,” *IEEE Transactions on Control of Network Systems*, vol. 6, no. 1, pp. 249–260, 2018.

[10] F. Rossi, R. Iglesias, M. Alizadeh, and M. Pavone, “On the interaction between Autonomous Mobility-on-Demand systems and the power network: Models and coordination algorithms,” *IEEE Transactions on Control of Network Systems*, vol. 7, no. 1, pp. 384–397, 2019.

[11] B. Turan and M. Alizadeh, “Competition in electric Autonomous Mobility-on-Demand systems,” *IEEE Transactions on Control of Network Systems*, to appear, 2021.

[12] B. Turan, R. Pedarsani, and M. Alizadeh, “Dynamic pricing and fleet management for electric autonomous mobility on demand systems,” *Transportation Research Part C: Emerging Technologies*, vol. 121, p. 102829, 2020.

[13] N. Tucker, B. Turan, and M. Alizadeh, “Online charge scheduling for electric vehicles in autonomous mobility on demand fleets,” in *IEEE Intelligent Transportation Systems Conference (ITSC)*, 2019, pp. 226–231.

[14] N. Tucker and M. Alizadeh, “An online admission control mechanism for electric vehicles at public parking infrastructures,” *IEEE Transactions on Smart Grid*, vol. 11, no. 1, pp. 161–170, 2019.

[15] A. Moradipari and M. Alizadeh, “Pricing and routing mechanisms for differentiated services in an electric vehicle public charging station network,” *IEEE Transactions on Smart Grid*, vol. 11, no. 2, pp. 1489–1499, 2019.

[16] S. Amani, M. Alizadeh, and C. Thrampoulidis, “Linear stochastic bandits under safety constraints,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2019, pp. 9256–9266.

[17] A. Moradipari, C. Thrampoulidis, and M. Alizadeh, “Stage-wise conservative linear bandits,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.

[18] A. Moradipari, S. Amani, M. Alizadeh, and C. Thrampoulidis, “Safe linear thompson sampling with side information,” *IEEE Transactions on Signal Processing*, to appear, 2021.

[19] S. Amani, M. Alizadeh, and C. Thrampoulidis, “Regret bounds for safe Gaussian process bandit optimization,” in *IEEE International Symposium on Information Theory (ISIT)*, 2021.

[20] N. Tucker, A. Moradipari, and M. Alizadeh, “Constrained Thompson sampling for real-time electricity pricing with grid reliability constraints,” *IEEE Transactions on Smart Grid*, vol. 11, no. 6, pp. 4971–4983, 2020.

[21] A. Moradipari, M. Ghavamzadeh, and M. Alizadeh, “Collaborative multi-agent stochastic linear bandits,” [Preprint URL].

[22] A. Moradipari, Y. Abbasi-Yadkori, M. Alizadeh, and M. Ghavamzadeh, "Parameter and feature selection in stochastic linear bandits," [Preprint URL].

[23] B. Turan, C. A. Uribe, H.-T. Wai, and M. Alizadeh, "Resilient primal-dual optimization algorithms for distributed resource allocation," *IEEE Transactions on Control of Network Systems*, to appear, 2020.

[24] B. Turan, C. Uribe, H. T. Wai, and Alizadeh, "On robustness of the normalized subgradient method with randomly corrupted subgradients," in *American Control Conference (ACC)*, 2021.

[25] B. Turan, C. Uribe, H. T. Wai, and Alizadeh, "On robustness of normalized block coordinate descent method for non-convex optimization," [Preprint URL].

[26] B. Turan, C. A. Uribe, H.-T. Wai, and M. Alizadeh, "Robust distributed optimization with randomly compromised gradients," [Preprint URL].

[27] K. Paarporn, B. Canty, P. N. Brown, M. Alizadeh, and J. R. Marden, "The impact of complex and informed adversarial behavior in graphical coordination games," *IEEE Transactions on Control of Network Systems*, vol. 8, no. 1, pp. 200–211, 2021.

[28] K. Paarporn, M. Alizadeh, and J. R. Marden, "A risk-security tradeoff in graphical coordination games," *IEEE Transactions on Automatic Control*, vol. 66, no. 5, pp. 1973–1985, 2021.

[29] K. Paarporn, R. Chandan, M. Alizadeh, and J. R. Marden, "The division of assets in multiagent systems: A case study in team blotto games," [Preprint URL].

[30] K. Paarporn, R. Chandan, M. Alizadeh, and J. R. Marden, "Characterizing the interplay between information and strength in Blotto games," in *IEEE 58th Conference on Decision and Control (CDC)*, 2019, pp. 5977–5982.

[31] K. Paarporn, R. Chandan, M. Alizadeh, and J. R. Marden, "A general lotto game with asymmetric budget uncertainty," [Preprint URL].