Resilience in Multi-Agent Systems: Recent Contributions

James Usevitch
Department of Aerospace Engineering
University of Michigan, Ann Arbor
Email: usevitch@umich.edu

Dimitra Panagou
Department of Aerospace Engineering
University of Michigan, Ann Arbor
Email: dpanagou@umich.edu

***Abstract*—The property of resilience describes the ability of a multi-agent system to mitigate the effects of adversarial attacks or faults. A growing body of resilient control techniques based upon the Mean-Subsequence-Reduced (MSR) class of algorithms have been proposed which enable the consensus of normally behaving agents despite attacks and faults. Impressive results have been obtained using these algorithms, but there is still much potential for future research. In this paper, several recent contributions of the authors to this area are summarized including results dealing with robustness determination, resilient leader-follower methods, and incorporating finite-time convergence into resilient control. We conclude with several opportunities for future work.**

I. Introduction

Resilience, or the ability to overcome the effects of faults and adversarial attacks, is an essential property for any multi-agent system intending to operate in the modern world. The literature is replete with examples of attacks that have been successfully carried out against systems in practice [1], [2]. However, many algorithms and control strategies do not account for the possibility of such attacks and faults. Incorporating resilience into multi-agent systems is currently an active field of research. A high level survey of network resilience and how it fits into the larger picture of security of cyber-physical systems is given in [1].

A growing body of literature has started incorporating resilience into multi-agents systems by using several variants of the *mean-subsequence-reduced* (MSR) family of algorithms [3]. Several representative examples of such algorithms can be found in [4]–[12]. In essence, these algorithms address the *resilient consensus* problem where normally behaving agents must come to agreement on information in the presence of a bounded number of adversarial agents with unknown identity. MSR-type algorithms have the advantage in that they are simple to implement and rely only on local data. However, the price of this simplicity are the requirements imposed on the network structure to guarantee the success of the algorithms. Many of the sufficient conditions for resilient consensus in these paper rely upon the graph theoretic notions of *r-robustness, strong r-robustness*, and *(r, s)-robustness* [4], [13]. The conditions of r-robustness and (r, s)-robustness are computationally hard to determine for general graphs [13], [14], which has been a key challenge in implementing MSR-type algorithms in practice.

In this paper we summarize several of our recent contributions to the resilient control literature. Our work can be organized into two general categories: techniques for analyzing network robustness, and theory which incorporates resilience into control strategies. Our recent work in relation to these two categories is described in sections II and III, respectively.

II. Analyzing Network Robustness

The MSR-type resilient algorithms discussed previously can only guarantee convergence of the normal agents in the presence of adversaries if the communication network is sufficiently robust. However, a key challenge in implementing these resilient algorithms is that determining the r- and (r, s)-robustness of arbitrary digraphs is an NP-hard problem in general [13], [14]. Several of our recent contributions have therefore focused on determining the robustness of networks either exactly or approximately.

For some classes of graphs, lower bounds on the robustness can be established using the mathematical properties of the graphs [13], [15]–[20]. In [21] and [22], we analyzed robustness properties of a class of networks called k-circulant digraphs, where k is an integer parameter which determines the structure of the network; the reader is referred to [21] for additional details. In [21] we demonstrated that for all k-circulant digraphs, the maximum integer r for which a given digraph is r-robust is lower bounded by $\lceil k/2 \rceil$. This result allows for such digraphs to be created with a desired robustness level and scaled to a nearly arbitrary number of nodes.[1] In [22], we demonstrated the conditions under which k-circulant digraphs are also *strongly r-robust* [4] with respect to a subset of nodes $S \in \mathcal{V}$.

However, most of the aforementioned methods are restricted in their scope because 1) they apply only to particular classes of graphs (e.g. undirected graphs, k-circulant graphs, triangular robust graphs, etc.), 2) many of the methods only provide fixed approximate lower bounds which cannot be iteratively tightened on the maximum integer r for which a given graph is r-robust, and 3) most do not consider the more general property of (r, s)-robustness. The $CheckRobustness$ and $DetermineRobustness$ algorithms in [14] are notable

[1] For a given value of k, a k-circulant digraph requires $n \geq k+1$ nodes. The number of nodes can be scaled *up* to any arbitrary value from this minimum number.

exceptions, which are able to analyze the *exact* r- and (r, s)-robustness of *digraphs* in general. More precisely, these algorithms are able to determine the maximum r for which a digraph is r-robust, and the maximum (r, s)-pair with respect to a lexicographic ordering for which the digraph is (r, s)-robust. However, these algorithms are essentially exhaustive search techniques, have exponential complexity in the number of nodes in the digraph, and are only able to iteratively tighten *upper* bounds on the integers r and s for which a given digraph is r- or (r, s)-robust.

In [23] and [24], we present a formulation for determining the *exact* r-and (r, s)-robustness of nonempty, nontrivial, simple digraphs and undirected graphs using mixed integer linear programming (MILP). To the best of our knowledge, these results are the first to use optimization techniques to determine the robustness of digraphs. Applying MILP methods to the robustness determination problem provides several advantages. First, expressing the robustness determination problem in MILP form allows for approximate *lower* bounds on a given digraph's r-robustness to be iteratively tightened using algorithms such as branch-and-bound. Lower bounds on the maximum value of s for which a given digraph is (r, s) robust (for a given nonnegative integer r) can also be iteratively tightened using the approach in [23], [24]. Prior algorithms are only able to tighten the upper bound on the maximum robustness for a given digraph or undirected graph. Second, this formulation enables commercially available solvers such as Gurobi or MATLAB's *intlinprog* to be used to find the maximum robustness of any digraph. Algorithmic advances and improvement in computer hardware have led to a speedup factor of 800 billion for mixed integer optimization problems during the last 25 years [25], and the results of [23], [24] allow for the robustness determination problem to benefit from these ongoing and future improvements.

III. Resilience in Control Strategies

Much of the prior literature incorporating MSR-type algorithms have focused on consensus to values within the convex hull of initial normal agent states at an asymptotic or exponential convergence rate. There are many possibilities to both extend the resilient characteristics of such algorithms to additional control objectives and incorporate alternate methods of filtering out adversarial information. In this section we outline we outline our recent work which explores both possibilities.

In [22], we extend the resilient characteristics of the W-MSR algorithm to a leader-follower scenario, where normal agents must track the state of a set of leaders while filtering out adversarial agents. The challenging aspect of the problem lies in the fact that normal agents do not know the identity of their in-neighbors, i.e. whether in-neighbors are normal, adversarial, or leaders. In addition, the leaders' state value may lie outside the convex hull of initial normal agents' states. The paper [22] first shows that the presence of at least $F+1$ leaders is a necessary condition for leader-follower consensus. It then also demonstrates that selecting an arbitrary set of $F+1$ agents in a $(2F+1)$-robust or $(F+1, F+1)$-robust graph is *not* sufficient to guarantee that the normal agents resiliently converge to the leader agents. Rather, we demonstrate in this paper that under an F-local adversary model, the normal agents converge to a set of $2F+1$ or more leader agents $\mathcal{L}$ if the network is *strongly* $(2F+1)$-*robust* with respect to $\mathcal{L}$, and the value of the leader agents remains constant. A useful aspect of this result is that convergence of the normal agents to the leader agents is guaranteed even if up to F of the *leader* agents become adversarial. The case of incorporating *trusted nodes* [26], [27] is also considered, where it is assumed that leader agents have been sufficiently secured as to render them immune to adversarial attacks. We introduce the notion of *trusted leader-follower robustness* in this case which is a sufficient condition for the convergence of normal agents to a set of at least $F+1$ leader agents in a network under an F-local adversarial model.

In [28] we incorporate several novel elements into a resilient control scenario. Works published prior to [28] almost exclusively consider algorithms with asymptotic or exponential convergence. A notable exception is [29], where resilient consensus is achieved in finite-time. However, [29] considers only undirected graphs under the assumption that all possibly misbehaving nodes are only connected to trusted nodes which are guaranteed to be cooperative. Our work in [28] introduces a novel continuous finite-time controller that allows agents to achieve formations in the presence of adversarial agents. As opposed to [29], we do not assume the existence of any trusted agents, but rather consider the more general F-total adversarial model. The controller employs a novel filtering mechanism based on the norm of the difference between agents' states. In addition, it is proven that this controller guarantees convergence with bounded inputs. To achieve this, we define novel conditions for the filtering timing and input weights which ensure that agents can remain in formation even with a dwell time in the filtering mechanism. Finally, we also show in [28] that the norm-based filtering and bounded input elements of our continuous-time controller can be used in a similar resilient discrete-time system, which is proven to have exponential convergence.

IV. Conclusion and Future Work

The need for multi-agent systems resilient to misinformation and misbehavior will only increase as distributed systems become more widespread and ubiquitous. Creating algorithms and control strategies guaranteeing such resilience continues to be a highly challenging problem, but simulaneously a highly interesting and important problem. Specifically for the area of network resilience, several potential research directions of immediate interest include the following:

Adaptation of resilience methods to mobile robotics: Creating multi-agent mobile robotic systems has inherent challenges which include asynchronous systems; limited communication ranges; noise in sensor measurements and communication; and limited power and computational resources. The excellent work in [6], [18]–[20], [30], [31] represents some of the first efforts to adapt MSR-type resilient algorithms to mobile agents with limited communication radii and time-varying networks. However, much work remains to be done in adapting these resilient algorithms to the challenges and control objectives of mobile robots. In particular, hardware implementations of such resilient algorithms on physical platforms are still relatively rare.

Extensions to systems with stochasticity and noise: The majority of papers in the resilience literature based upon MSR-type algorithms deal with deterministic systems. Notable exceptions incorporating stochasticity include [10], [32], where randomness is intentionally introduced by design into the agents' control laws to ensure convergence, and [33] where the effects of random packet drops are considered in the analysis of a resilient estimation algorithm. In physical systems, the information available to agents inevitably includes uncertainty and noise. Additional theoretical work is needed to evaluate the efficacy of MSR-type algorithms in such stochastic scenarios.

More intelligent identification of faulty and adversarial information: Agents applying MSR-type algorithms are able to filter out the effects of adversarial information by using a very simple filtering mechanism. An interesting and challenging direction for future research is developing more sophisticated methods for agents to identify adversarial information using only the local data available to them. Leveraging data-based or learning-based techniques towards this end might enable the strict requirements on the network communication topologies of resilient networks to be relaxed.

References

[1] S. M. Dibaji, M. Pirani, D. B. Flamholz, A. M. Annaswamy, K. H. Johansson, and A. Chakrabortty, "A systems and control perspective of cps security," 2019.

[2] Y. Cherdantseva, P. Burnap, A. Blyth, P. Eden, K. Jones, H. Soulsby, and K. Stoddart, "A review of cyber security risk assessment methods for scada systems," *Computers & security*, vol. 56, pp. 1–27, 2016.

[3] R. M. Kieckhafer and M. H. Azadmanesh, "Reaching approximate agreement with mixed-mode faults," *IEEE Transactions on Parallel and Distributed Systems*, vol. 5, no. 1, pp. 53–63, 1994.

[4] H. Zhang and S. Sundaram, "Robustness of information diffusion algorithms to locally bounded adversaries," in *American Control Conference (ACC), 2012*, pp. 5855–5861, IEEE, 2012.

[5] H. J. LeBlanc, H. Zhang, X. Koutsoukos, and S. Sundaram, "Resilient asymptotic consensus in robust networks," *IEEE Journal on Selected Areas in Communications*, vol. 31, no. 4, pp. 766–781, 2013.

[6] D. Saldana, A. Prorok, S. Sundaram, M. F. Campos, and V. Kumar, "Resilient consensus for time-varying networks of dynamic agents," in *American Control Conference (ACC), 2017*, pp. 252–258, IEEE, 2017.

[7] A. Mitra and S. Sundaram, "Secure distributed state estimation of an lti system over time-varying networks and analog erasure channels," in *2018 Annual American Control Conference (ACC)*, pp. 6578–6583, June 2018.

[8] A. Mitra and S. Sundaram, "Secure distributed observers for a class of linear time invariant systems in the presence of byzantine adversaries," in *Decision and Control (CDC), 2016 IEEE 55th Conference on*, pp. 2709–2714, IEEE, 2016.

[9] S. M. Dibaji and H. Ishii, "Resilient consensus of second-order agent networks: Asynchronous update rules with delays," *Automatica*, vol. 81, pp. 123–132, 2017.

[10] S. M. Dibaji, H. Ishii, and R. Tempo, "Resilient randomized quantized consensus," *IEEE Transactions on Automatic Control*, vol. 63, no. 8, pp. 2508–2522, 2018.

[11] H. J. LeBlanc and X. Koutsoukos, "Resilient first-order consensus and weakly stable, higher order synchronization of continuous-time networked multi-agent systems," *IEEE Transactions on Control of Network Systems*, 2017.

[12] H. J. LeBlanc, H. Zhang, S. Sundaram, and X. Koutsoukos, "Resilient continuous-time consensus in fractional robust networks," in *2013 American Control Conference*, pp. 1237–1242, 6 2013.

[13] H. Zhang, E. Fata, and S. Sundaram, "A notion of robustness in complex networks," *IEEE Transactions on Control of Network Systems*, vol. 2, no. 3, pp. 310–320, 2015.

[14] H. J. LeBlanc and X. D. Koutsoukos, "Algorithms for determining network robustness," in *Proceedings of the 2nd ACM international conference on High confidence networked systems*, pp. 57–64, ACM, 2013.

[15] E. M. Shahrivar, M. Pirani, and S. Sundaram, "Spectral and structural properties of random interdependent networks," *Automatica*, vol. 83, pp. 234–242, 2017.

[16] J. Zhao, O. Yağan, and V. Gligor, "On connectivity and robustness in random intersection graphs," *IEEE Transactions on Automatic Control*, vol. 62, no. 5, pp. 2121–2136, 2017.

[17] E. M. Shahrivar, M. Pirani, and S. Sundaram, "Robustness and algebraic connectivity of random interdependent networks," *arXiv preprint arXiv:1508.03650*, 2015.

[18] D. Saldana, A. Prorok, M. F. Campos, and V. Kumar, "Triangular networks for resilient formations," in *Distributed Autonomous Robotic Systems*, pp. 147–159, Springer, 2018.

[19] L. Guerrero-Bonilla, D. Saldana, and V. Kumar, "Design guarantees for resilient robot formations on lattices," *IEEE Robotics and Automation Letters*, vol. 4, no. 1, pp. 89–96, 2019.

[20] D. Saldaña, L. Guerrero-Bonilla, and V. Kumar, "Resilient backbones in hexagonal robot formations," in *Distributed Autonomous Robotic Systems*, pp. 427–440, Springer, 2019.

[21] J. Usevitch and D. Panagou, "r-robustness and (r, s)-robustness of circulant graphs," in *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, pp. 4416–4421, Dec 2017.

[22] J. Usevitch and D. Panagou, "Resilient leader-follower consensus to arbitrary reference values," in *2018 Annual American Control Conference (ACC)*, pp. 1292–1298, June 2018.

[23] J. Usevitch and D. Panagou, "Determining r-robustness of arbitrary digraphs using zero-one linear integer programming," in *2019 American Control Conference, to appear*.

[24] J. Usevitch and D. Panagou, "Determining r-and (r, s)-robustness of digraphs using mixed integer linear programming," *arXiv preprint arXiv:1901.11000*, 2019.

[25] D. Bertsimas and J. Dunn, "Optimal classification trees," *Machine Learning*, vol. 106, no. 7, pp. 1039–1082, 2017.

[26] W. Abbas, A. Laszka, and X. Koutsoukos, "Improving network connectivity and robustness using trusted nodes with application to resilient consensus," *IEEE Transactions on Control of Network Systems*, vol. 5, no. 4, pp. 2036–2048, 2018.

[27] A. Mitra, W. Abbas, and S. Sundaram, "On the impact of trusted nodes in resilient distributed state estimation of lti systems," in *2018 IEEE Conference on Decision and Control (CDC)*, pp. 4547–4552, IEEE, 2018.

[28] J. Usevitch, K. Garg, and D. Panagou, "Finite-time resilient formation control with bounded inputs," in *2018 IEEE Conference on Decision and Control (CDC)*, pp. 2567–2574, Dec 2018.

[29] M. Franceschelli, A. Giua, and A. Pisano, "Finite-time consensus on the median value with robustness properties," *IEEE Transactions on Automatic Control*, vol. 62, no. 4, pp. 1652–1667, 2017.

[30] L. Guerrero-Bonilla, A. Prorok, and V. Kumar, "Formations for Resilient Robot Teams," in *IEEE Robotics and Automation Letters*, vol. 2, pp. 841–848, IEEE, 2017.

[31] K. Saulnier, D. Saldana, A. Prorok, G. J. Pappas, and V. Kumar, "Resilient flocking for mobile robot teams," *IEEE Robotics and Automation Letters*, vol. 2, no. 2, pp. 1039–1046, 2017.

[32] M. Nakamura, H. Ishii, and S. M. Dibaji, "Maximum-based consensus and its resiliency," *IFAC-PapersOnLine*, vol. 51, no. 23, pp. 283–288, 2018.

[33] A. Mitra, J. A. Richards, S. Bagchi, and S. Sundaram, "Resilient distributed state estimation with mobile agents: overcoming byzantine adversaries, communication losses, and intermittent measurements," *Autonomous Robots*, pp. 1–26, 2018.