

Stochastic and Deterministic Fault Detection for Randomized Gossip Algorithms

Daniel Silvestre ^{a,b}, Paulo Rosa ^{c,b}, João P. Hespanha ^d Carlos Silvestre ^{a,b}

^a*Department of Electrical and Computer Department, Faculty of Science and Technology, University of Macau, China.*

^b*Institute for Systems and Robotics, Instituto Superior Tecnico, Universidade de Lisboa, Lisboa, Portugal.*

^c*Deimos Engenharia, Lisbon, Portugal.*

^d*Department of Electrical and Computer Department, University of California, Santa Barbara, CA 93106-9560, USA.*

Abstract

This paper addresses the problem of detecting faults in linear randomized gossip algorithms, where the selection of the dynamics matrix is stochastic. A fault is a disturbance signal injected by an attacker to corrupt the states of the nodes. We propose the use of Set-Valued Observers (SVOs) to detect if the state observations are compatible with the system dynamics for the worst case in a deterministic setting. The concept of Stochastic Set-Valued Observers (SSVOs) is also introduced to construct a set that is guaranteed to contain all possible states with, at least, a pre-specified desired probability. The proposed algorithm is stable in the sense that it requires a finite number of vertices to represent polytopic sets and it allows for the computation of the largest magnitude of the disturbance that an attacker can inject in the network without being detected. Results are presented to reduce the computational cost of this approach and, in particular, by considering only local information and representing the remainder of the network as a disturbance. The case of a consensus algorithm is discussed leading to the conclusion that, by using the proposed SVOs, finite-time consensus is achieved in non-faulty environments. A novel algorithm is proposed that produces less conservative set-valued state estimates by having nodes exchanging local estimates. The algorithm inherits all the previous properties and also enables finite-time consensus computation regardless of the value of the horizon.

Key words: Fault detection, computer networks, decentralization, estimation theory, randomized methods, linear parameterically varying (LPV) methodologies.

1 Introduction

The problem of detecting faults in an asynchronous distributed environment relates to determining if any of the nodes enters in an incoherent state given the observed history of measurements. In particular, we are interested in randomized algorithms where the dynamics is common to all the nodes and no control messages are needed. This class of algorithms is used for iterative solutions because they offer a certain level of robustness against packet drops and node failure. Applications of randomized algorithms range from computing integrals to consensus [4] and solving problems for which the solution

requires a heavy computational burden [6], [10]. Large scale distributed systems and robot swarms highlight the importance of this problem for practical applications.

The aim of this paper is to detect the presence of an attacker who corrupts the states of the nodes or their transmissions. In this context, the small probability of an event cannot be discarded as an attacker can select the worst case signal, which motivates the use of set-valued estimation tools. Therefore, we address the problem in a distributed manner where each node models the network from its perspective as a Linear Parameter-Varying (LPV) system, where the input is the attacker signal. Since an attacker is allowed to inject any signal, we are looking at the worst case scenario and estimating the set of all possible state realizations that comply with the “fault-free” model. If the set becomes empty, we can guarantee the presence of an attacker or any other fault.

Fault detection methods have been proposed in the lit-

Email address: Contact the author D. Silvestre through dsilvestre@isr.ist.utl.pt for a free unformatted copy of this article or go to <https://www.sciencedirect.com/science/article/pii/S0005109816305192> for the final formatted version. (Daniel Silvestre).

erature for a number of specific applications. For instance, [7] focuses on defining multiple classes of theoretic unreliable fault detectors for the case of a consensus problem. The proposed method checks if the algorithm is running correctly and if all the messages are in concordance with the specifications. The work [11] considers the problem of detecting and correcting the state of the system in the presence of a fault, and make the distinction between a general fault or one produced by a malicious agent. Bounds on the number of corrupted nodes to ensure detectability are given for both cases. In [11], the system dynamics are described by a linear time-invariant model that constrains the network topology to be fixed and communications to be synchronous. Here, however, these assumptions are dropped as we consider a randomized gossip algorithm.

The adopted strategy for fault detection has an interesting finite-time property that can be used in consensus problems. Finite-time consensus in the presence of malicious agents has been addressed in [18], where the authors show that the topology of the network categorizes its ability to deal with attacks. Both the number of corrupted nodes and vertex-disjoint paths in the network influence its resilience. In [18], it is assumed a broadcast model where, at each transmission time, the nodes send to all their neighbors the same value and the agents objective is to compute some function of the initial states. The main difference to the work described herein is the communications model, which we assume to be *gossip*, where pairs of nodes are selected randomly to exchange information, instead of having a broadcast model.

A subset of the results described herein was previously presented in the conference papers [16] and [17] by the same authors. In [16], the concept of Stochastic Set-Valued Observers (SSVOs) was presented using α -confidence sets, i.e., sets where the state of the system is guaranteed to belong with a desired pre-specified $1 - \alpha$ probability; which can be viewed as a generalization of confidence intervals. The property of finite-time consensus when using (deterministic) Set-Valued Observers (SVOs) for a sufficiently large horizon in a randomized gossip consensus algorithm is shown in [17].

The main contributions of this paper are as follows:

- it is shown how to compute a threshold for the “maximum impact” of an undetected fault, discussing two particular cases: linear consensus, and networked physical systems;
- the number of required communications for guaranteeing detection is reduced by analyzing the structure of randomized gossip algorithms;
- finally, we show how some of the dynamics matrices can be discarded from the model that each node has of the network, which reduces the computational complexity of the fault detection procedure.

In this paper, we resort to the concept of SVOs for keep-

ing a polytopic state set-valued estimate, firstly introduced in [19] and [14]. For the interested reader, further information can be found in [15] and [9] and the references therein. An alternative to SVOs could be the use of zonotopes, described in [2] and further developed in [1]. Zonotopes represent a different trade-off between the computation complexity of unions and intersections. In particular, intersections introduce conservatism which motivated the alternative approach adopted in this article in order to attain the desired convergence guarantees, while keeping the tractable computational requirements.

The applicability of the proposed method in the detection of faults in randomized gossip algorithms spans other purposes as several challenges in the Fault Detection and Isolation (FDI) literature - [3] - share the framework described in the sequel. In [13], the authors take advantage of SVOs for fault detection by resorting to a model falsification approach. This paper extends the results in [13] to detect faults in randomized gossip algorithms by rewriting the associated dynamics as an LPV model. Moreover, unlike the approach in [13], the method proposed herein takes into account the information related to the probability of having a given communication, in order to reduce the conservatism of the results.

Closely related to the concept of stochastic detection is the work presented in [12] which performs the detection by finding the change points in the correlation statistics for a sensing network. The authors are able to provide guarantees on detection delay and false alarm probability. Such approach addresses a similar problem of detecting faults that are possible in the standard dynamics but not very “probable” to take place. Our work tackles this issue in a different way by considering the set of possible states given the more “probable” dynamics.

In the context of fault detection in distributed systems, [8] addresses the problem by looking at the whole system and constructing a batch of observers for each sub-system. By looking at the outputs of these observers it is possible to detect and isolate faults affecting one of the sub-systems. However, it is a centralized approach whereas our focus is to run each of the observers locally at each sub-system in a fully distributed way.

In [20], the authors propose an on-line fault detection and isolation algorithm for linear discrete-time uncertain systems where the detection is based on the computation of upper and lower bounds for the fault signal. The calculations are performed resorting to Linear Matrix Inequality (LMI) optimization techniques. Similar computational burden considerations to the work presented in this article are discussed and the techniques are related to our work. However, in order to address randomized gossip algorithms we studied a more general class of systems.

Using the approach of design residual filters, [5] studies a class of linear continuous-time systems with the purpose of identifying faulty actuators. The aim of this work is to adjust the filters parameters as to decouple them when faults affect a group of actuators. Our approach differs in the sense that we want to incorporate unknown parameters in the dynamics matrix of the system.

The organization of this paper develops towards presenting all the details of fault detection in the worst-case and stochastic senses for distributed linear systems. Initial focus is given to gossip systems and key constraints posed on the detection, namely, the characteristics associated with the network component and how faults are modeled. The concept of SVOs is introduced and applied to the deterministic fault detection, as the worst-case is considered. Progress is made in presenting a method to extend the SVOs computation to incorporate the stochastic information of the communication process, which results in the SSVOs.

The SVO-based fault detection method motivates the introduction of a consensus algorithm that performs averages on intervals containing the state, intersecting them upon neighbor communication. The algorithm is asymptotically convergent and, under some communication patterns, it finds the consensus value in finite-time due to the intersection phase. Therefore, this paper is proposing an SVO-based approach to fault detection with different types of SVOs. The deterministic worst-case detection is performed by an SVO running at each node using only locally available information. The stochastic detection is an extension of the previous method with the set of state estimates being a subset of the previous one corresponding to a confidence set of where the state can take values. Lastly, in the particular case of consensus, it is introduced an algorithm that takes advantage of the local estimates and intersects them upon communication to generate less conservative sets.

References

- [1] T. Alamo, J.M. Bravo, and E.F. Camacho. Guaranteed state estimation by zonotopes. *Automatica*, 41(6):1035 – 1043, 2005.
- [2] D. Bertsekas and I. Rhodes. Recursive state estimation for a set-membership description of uncertainty. *IEEE Transactions on Automatic Control*, 16(2):117 – 128, apr 1971.
- [3] J. Bokor and Z. Szabó. Fault detection and isolation in nonlinear systems. In *Annual Reviews in Control 33.2*, pages 113–123, 2009.
- [4] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah. Randomized gossip algorithms. *IEEE Transactions on Information Theory*, 52(6):2508 – 2530, June 2006.
- [5] A. Cristofaro and T.A. Johansen. Fault tolerant control allocation using unknown input observers. *Automatica*, 50(7):1891 – 1897, 2014.
- [6] M. Dietzfelbinger, T. Hagerup, J. Katajainen, and M. Penttonen. A reliable randomized algorithm for the closest-pair problem. *Journal of Algorithms*, 25(1):19 – 51, 1997.
- [7] K.P. Kihlstrom, L.E. Moser, and P.M. Melliar-Smith. Solving consensus in a byzantine environment using an unreliable fault detector. In *Proceedings of the International Conference on Principles of Distributed Systems (OPODIS)*, pages 61–75, 1997.
- [8] P.P. Menon and C. Edwards. Robust fault estimation using relative information in linear multi-agent networks. *IEEE Transactions on Automatic Control*, 59(2):477–482, Feb 2014.
- [9] M. Milanese and A. Vicino. Optimal estimation theory for dynamic systems with set membership uncertainty: An overview. *Automatica*, 27(6):997 – 1009, 1991.
- [10] K. Mulmuley. *Computational Geometry: An Introduction through Randomized Algorithms*. Prentice-Hall, NJ, 1994.
- [11] F. Pasqualetti, A. Bicchi, and F. Bullo. Consensus computation in unreliable networks: A system theoretic approach. *IEEE Transactions on Automatic Control*, 57(1):90 – 104, jan. 2012.
- [12] R. Rajagopal, X. Nguyen, S.C. Ergen, and P. Varaiya. Distributed online simultaneous fault detection for multiple sensors. In *International Conference on Information Processing in Sensor Networks, IPSN '08*, pages 133–144, April 2008.
- [13] P. Rosa and C. Silvestre. Fault detection and isolation of LPV systems using set-valued observers: An application to a fixed-wing aircraft. *Control Engineering Practice*, 21(3):242 – 252, 2013.
- [14] F. Schweppe. Recursive state estimation: Unknown but bounded errors and system inputs. *IEEE Transactions on Automatic Control*, 13(1):22 – 28, feb 1968.
- [15] F. Schweppe. *Uncertain Dynamic Systems*. Prentice-Hall, 1973.
- [16] D. Silvestre, P. Rosa, R. Cunha, J. P. Hespanha, and C. Silvestre. Gossip average consensus in a byzantine environment using stochastic set-valued observers. In *52nd IEEE Conference on Decision and Control*, 2013.
- [17] D. Silvestre, P. Rosa, J. P. Hespanha, and C. Silvestre. *Finite-time Average Consensus in a Byzantine Environment Using Set-Valued Observers*. American Control Conference, 2014.
- [18] S. Sundaram and C.N. Hadjicostis. Distributed function calculation via linear iterative strategies in the presence of malicious agents. *IEEE Transactions on Automatic Control*, 56(7):1495–1508, July 2011.
- [19] H. Witsenhausen. Sets of possible states of linear systems given perturbed observations. *IEEE Transactions on Automatic Control*, 13(5):556 – 558, oct 1968.
- [20] Z. Zhang and I.M. Jaimoukha. On-line fault detection and isolation for linear discrete-time uncertain systems. *Automatica*, 50(2):513 – 518, 2014.