Secure state-estimation for dynamical systems under active adversaries

Hamza Fawzi, Paulo Tabuada and Suhas Diggavi
Department of Electrical Engineering,
University of California at Los Angeles, Los Angeles, CA 90095-1594
{hfawzi,tabuada,suhas}@ee.ucla.edu

***Abstract*—We consider the problem of state-estimation of a linear dynamical system when some of the sensor measurements are corrupted by an adversarial attacker. The errors injected by the attacker in the sensor measurements can be arbitrary and are not assumed to follow a specific model (in particular they can be of arbitrary magnitude). We first characterize the number of attacked sensors that can be tolerated so that the state of the system can still be correctly recovered by any decoding algorithm. We then propose a specific computationally feasible decoding algorithm and we give a characterization of the number of errors this decoder can correct. For this we use ideas from compressed sensing and error correction over the reals and we exploit the dynamical nature of the problem. We show using numerical simulations that this decoder performs very well in practice and allows to correct a large number of errors.**

I. Introduction

Everyday control systems work silently in the background to support much of the critical infrastructure we have grown used to. Water distribution networks, sewer networks, gas and oil networks, and the power grid are just a few examples of critical infrastructure that rely on control systems for its normal operation. These systems are becoming increasingly networked both for distributed control and sensing, as well as for remote monitoring and reconfiguration. Unfortunately, once these systems become connected to the internet they become vulnerable to attacks that, although launched in the cyber domain, have as objective to manipulate the physical part. This realization led to the emergence of new security challenges that are distinct from traditional cyber security as highlighted in [1]–[3]. The importance of these questions cannot be understated, and has been recognized by several federal agencies including the Department of Homeland Security [2], [3].

The design of control systems that work correctly under faults and failures is certainly not a new problem. Existing design methodologies include fault-tolerant control [4], robust control [5], game theoretic techniques [6], and error-correction. All of these approaches are based on certain assumptions on the nature of the faults, failures, or disturbances. For example in fault-tolerant control it is usually assumed that there is a finite number of failure modes and that each mode follows a specific failure model. In robust control the disturbances are generally bounded and localized, i.e., they enter the plant model in a specific location. In game theoretic methods, the adversary objective, i.e., its utility function, is assumed to be known. Finally, in the field of error correction the channel noise is generally stochastic and its probabilistic model is also usually considered to be known by the decoder.

Since none of these assumptions are fully justifiable in the context of an adversarial attack, there has been a recent increase in control systems security research [7]–[11]. In the paper [7], the authors consider the problem of control and estimation in a networked system when the communication links are subject to disturbances. The disturbances are however assumed to follow a certain stochastic process which does not necessarily capture the behavior of an attacker. In [8] the authors consider a more intelligent jammer who plans his attacks in order to maximize a certain objective. This is however a strong assumption since the objective of the attacker is generally not known. In [10], [12] the problem of reaching consensus in the presence of malicious agents is considered. There, the authors give a characterization of the number of infected nodes that can be tolerated and, when possible, propose a way to detect and overcome the effect of the malicious agents. The methods that are proposed are however computationally expensive and are mainly of combinatorial nature (in [12] though, the structure of the consensus network is exploited and is used to design a more computationally efficient method). Moreover, in these works, the dynamics is part of the consensus algorithm and can be specifically designed, rather than being given, as in a physical system. Finally, there has also been recent work in the area of error-correction over adversarial channels, e.g., [13]. However the dynamics of the system does not generally play a role and the correction capability is studied in a static setting that does not take advantage of the dynamics of the system.

In this paper we consider the problem of robust state estimation of a discrete-time linear plant when some of the sensors are attacked. Unlike most of the works we mentioned above, we do not restrict the type of errors introduced by the attacker on the captured sensors (in particular the errors injected can be of arbitrary magnitude). We characterize the resilience in terms of the maximum number of attacked sensors that can be tolerated for correct state estimation. Our formulation of the problem can be seen as an extension of error correction where we take advantage of the known dynamics of the system, that

This work was partially supported by NSF awards 0820061 and 1136174.

acts like a natural error-correcting code. In fact we will see throughout the paper that the static error correction problem is a special case of the problem we consider when the decoding is done at every time step, without delay. We also propose a computationally feasible decoding algorithm to recover the state despite the errors. Our approach for the computationally efficient schemes uses ideas from the area of compressed sensing and its application in the context of error correction over the reals [13].

The paper is organized as follows. In section III we characterize, for a given system, the number of errors that can be corrected (by any decoder) and we show in particular that this number is maximal for "almost all" linear systems. We then propose in section IV an optimization-based decoder that can correct the maximum number of errors for a given linear plant. We show however that this decoder is not computationally feasible and, using ideas from compressed sensing, we propose in section V a relaxed decoder that is computationally feasible and that takes the form of a convex program. We also characterize the number of errors that can be corrected by this decoder. In section VI we demonstrate the performance of this decoder on numerical examples.

II. Notation and problem formulation

A. Notation

We use the following notations throughout the paper. If $S \subset S'$ is a set, we denote by $|S|$ its cardinality and by $S^c = S' \backslash S$ its complement (the parent set S' will be clear from the context). For a vector $x \in \mathbb{R}^n$, the support of x, denoted by $\mathsf{supp}(x)$, is the set of nonzero components of x:

$$\mathsf{supp}(x) = \{i \in \{1, \dots, n\} \mid x_i \neq 0\}.$$

The number of nonzero components of x will be denoted by $\|x\|_{\ell_0}$:

$$\|x\|_{\ell_0} = |\mathsf{supp}(x)|.$$

Also, if $K \subset \{1, \dots, n\}$, we let $\mathcal{P}_K$ be the projection map onto the components of K ($\mathcal{P}_K x$ is a vector with $|K|$ components).

For a matrix $M \in \mathbb{R}^{m \times n}$ we denote by $M_i \in \mathbb{R}^n$ the i'th row of M, for $i \in \{1, \dots, m\}$. We define the *row support* of M to be the set of nonzero rows of M:

$$\mathsf{rowsupp}(M) = \{i \in \{1, \dots, m\} \mid M_i \neq 0_{\mathbb{R}^n}\}.$$

As with the vector case, the notation $\|M\|_{\ell_0}$ will denote the number of nonzero rows of M:

$$\|M\|_{\ell_0} = |\mathsf{rowsupp}(M)|.$$

B. Problem formulation

In this section we define the problem to be solved in this paper. We also explicitly state all the assumptions we make regarding the plant, the communication, and the attacker.

1) Plant model: Consider a linear dynamical system given by

$$\begin{aligned} x^{(t+1)} &= Ax^{(t)} \\ y^{(t)} &= Cx^{(t)} + e^{(t)} \end{aligned} \tag{1}$$

where $x^{(t)} \in \mathbb{R}^n$ is the state at time $t \in \mathbb{N}$, $A \in \mathbb{R}^{n \times n}$ is the system matrix, $C \in \mathbb{R}^{p \times n}$ is the sensors measurement matrix, and $e^{(t)} \in \mathbb{R}^p$ are the errors injected by the malicious agent. If sensor $i \in \{1, \dots, p\}$ is not attacked then necessarily $e_i^{(t)} = 0$ and the output $y_i^{(t)}$ of sensor i is not corrupted, otherwise $e_i^{(t)}$ (and therefore $y_i^{(t)}$) can take any value. The sparsity pattern of the errors $e^{(t)}$ therefore indicates the set of attacked sensors.

2) Attacker model: We will assume in this paper that the set of attacked sensors does not change over time. More precisely, if $K \subset \{1, \dots, p\}$ is the set of sensors that were attacked, then we have for all t, $\mathsf{supp}(e^{(t)}) \subset K$. Note that this is a valid and realistic assumption when the time it takes for the malicious agent to gain control of a sensor is large compared to the time scale of the system. Furthermore observe that a model where the set of attacked sensors is allowed to change at every time step while having a fixed cardinality would in turn not be very realistic since it would assume that the attacker abandons from t to $t+1$ some of the sensors he had control over. For these reasons, we will assume for our model that the set K of attacked sensors is constant over time (and, of course, unknown).

Moreover, since we are dealing with a malicious agent, we will not assume the errors $e_i^{(t)}$ (for an attacked sensor i) to follow any particular model and we will simply take them to be arbitrary real numbers. The only assumption concerning the malicious agent will be about the *number* of sensors that were attacked. Our statements will then typically characterize the number of errors that can be corrected by a decoder, i.e., the number of attacked sensors that we can tolerate so that we can still unambiguously recover the correct state $x^{(0)}$.

3) Communication model: At every time step t the output $y_i^{(t)}$ of each sensor $i \in \{1, \dots, p\}$ is transmitted to some device, such as a controller or a monitoring station, whose objective is to reconstruct the state of the plant (we assume the transmission to be exact, i.e., noiseless, and we also assume that the device knows the matrices A and C in order to reconstruct the state). Note that since the device knows the dynamics of the system, the problems of reconstructing the current state $x^{(t)}$ or the initial condition $x^{(0)}$ are (at least theoretically, and when A is invertible) equivalent. In this paper we will therefore focus on the problem of reconstructing the initial state $x^{(0)}$.

III. Correction in finite time

Let $x^{(0)} \in \mathbb{R}^n$ be the initial state of the plant and let $y^{(0)}, \dots, y^{(T-1)} \in \mathbb{R}^p$ be the vectors that are transmitted from the sensors to the receiver device in the first T time steps. As we saw earlier, these vectors are given by

$$y^{(t)} = CA^t x^{(0)} + e^{(t)},$$

where $e^{(t)}$ represent the errors injected by the attacker. In fact, we have $\mathsf{supp}(e^{(t)}) \subset K$ with $K \subset \{1,\dots,p\}$ being the set of sensors that were attacked and whose data is unreliable.

Having received the T vectors $y^{(0)},\dots,y^{(T-1)}$, the receiver uses a decoder $D\colon (\mathbb{R}^p)^T \to \mathbb{R}^n$ in order to estimate the initial state $x^{(0)}$ of the plant. The decoder correctly estimates the initial state if $D(y^{(0)},\dots,y^{(T-1)}) = x^{(0)}$.

We will say that the decoder D corrects q errors if it correctly recovers the initial state $x^{(0)}$ for any set K of attacked sensors of cardinality less than or equal to q. More formally we introduce the following definition:

Definition 1. We say that q errors are correctable after T steps by the decoder $D\colon (\mathbb{R}^p)^T \to \mathbb{R}^n$ if for any $x^{(0)} \in \mathbb{R}^n$, and for any sequence of vectors $e^{(0)},\dots,e^{(T-1)}$ in $\mathbb{R}^p$ such that $\mathsf{supp}(e^{(t)}) \subset K$ with $|K| = q$, we have $D(y^{(0)},\dots,y^{(T-1)}) = x^{(0)}$ where $y^{(t)} = CA^t x^{(0)} + e^{(t)}$, $t = 0,\dots,T-1$.
Furthermore, we will say that q errors are correctable after T steps if there exists a decoder that can correct q errors after T steps.

Let $E_{q,T}$ denote the set of error vectors $(e^{(0)},\dots,e^{(T-1)}) \in (\mathbb{R}^p)^T$ that satisfy $\forall t \in \{0,\dots,T-1\}$, $\mathsf{supp}(e^{(t)}) \subset K$ for some $K \subset \{1,\dots,p\}$ with $|K| = q$. Note that $E_{q,T}$ is a union of $\binom{p}{q}$ subspaces in $(\mathbb{R}^p)^T$.

Observe that, by definition 1, the existence of a decoder that can correct q errors is equivalent to saying that the following map

$$\begin{aligned} \mathbb{R}^n \times E_{q,T} &\to (\mathbb{R}^p)^T \\ (x^{(0)}, e^{(0)},\dots,e^{(T-1)}) &\mapsto (y^{(0)},\dots,y^{(T-1)}) \\ &\quad = (Cx^{(0)} + e^{(0)},\dots, \\ &\qquad CA^{T-1}x^{(0)} + e^{(T-1)}) \end{aligned} \tag{2}$$

is invertible, or, more precisely, that it has an inverse for the first n components of its domain (we are only interested in the state $x^{(0)}$, and not necessarily the error vectors). However it is easy to see that these two conditions are equivalent since the error vectors are uniquely determined by $x^{(0)}$ and the $y^{(t)}$'s and are given by $e^{(t)} = y^{(t)} - CA^t x^{(0)}$. Thus expressing injectivity of this map is equivalent to saying that q errors are correctable. This gives the following proposition:

Proposition 1. Let $T \in \mathbb{N}\backslash\{0\}$. The following are equivalent: (i) There is *no* decoder that can correct q errors after T steps; (ii) There exists $x_a, x_b \in \mathbb{R}^n$ with $x_a \neq x_b$, and error vectors $(e_a^{(0)},\dots,e_a^{(T-1)}) \in E_{q,T}$ and $(e_b^{(0)},\dots,e_b^{(T-1)}) \in E_{q,T}$ such that $A^t x_a + e_a^{(t)} = A^t x_b + e_b^{(t)}$ for all $t \in \{0,\dots,T-1\}$.

The proposition above simply says that it is not possible to unambiguously recover the state $x^{(0)}$ if there are two distinct values x_a and x_b with $x_a \neq x_b$ that can, with less than q corrupted sensors, explain the received data.

Note that the domain of the map defined in (2) is the Cartesian product of the whole $\mathbb{R}^n$ with the error set $E_{q,T}$ which is unbounded. This means that we require the decoder to recover *any* initial state $x^{(0)}$ for *any* sequence of error vectors from $E_{q,T}$. In practice however one could consider only vectors $x^{(0)}$ in some set $\Omega \subset \mathbb{R}^n$ if one has prior knowledge on the initial state (for example, if the states are all nonnegative, say for physical reasons, then one could take $\Omega = \mathbb{R}^n_+$). Similarly, if the attacker has a finite amount of energy then we could envisage considering only elements of $E_{q,T}$ in a certain ball of finite radius. We do not however pursue this here, and we assume in particular that the initial state of the plant can be anywhere in $\mathbb{R}^n$ and that the magnitude of the errors can be arbitrary.

We now give a simple necessary and sufficient condition in terms of the matrices A and C for q errors to be correctable.

Proposition 2. Let $T \in \mathbb{N}\backslash\{0\}$. The following are equivalent: (i) There is a decoder that can correct q errors after T steps; (ii) For all $z \in \mathbb{R}^n\backslash\{0\}$, $|\mathsf{supp}(Cz) \cup \mathsf{supp}(CAz) \cup \cdots \cup \mathsf{supp}(CA^{T-1}z)| > 2q$.

Proof: **(i)**$\Rightarrow$ **(ii)**: Suppose for the sake of contradiction that there exists $z \in \mathbb{R}^n\backslash\{0\}$ such that $|\mathsf{supp}(Cz) \cup \mathsf{supp}(CAz) \cup \cdots \cup \mathsf{supp}(CA^{T-1}z)| \leq 2q$. Let $e_a^{(t)}$ and $e_b^{(t)}$ be such that $CA^t z = e_a^{(t)} - e_b^{(t)}$ with $\mathsf{supp}(e_a^{(t)}) \subset L_a$ and $\mathsf{supp}(e_b^{(t)}) \subset L_b$ with $|L_a| \leq q$ and $|L_b| \leq q$ (L_a and L_b are any two subsets of $\{1,\dots,p\}$ with cardinality less than or equal to q that satisfy $L_a \cup L_b = \mathsf{supp}(Cz) \cup \cdots \cup \mathsf{supp}(CA^{T-1}z)$). Now let, for $t \in \{0,\dots,T-1\}$, $y^{(t)} = CA^t z + e_b^{(t)} = CA^t \cdot 0 + e_a^{(t)}$. If q errors were correctable after T steps by some decoder D then we would have $D(y^{(0)},\dots,y^{(T-1)}) = z$ and also $D(y^{(0)},\dots,y^{(T-1)}) = 0$ which is impossible since $z \neq 0$.

(ii)$\Rightarrow$ **(i)**: We again resort to contradiction. Suppose that q errors are *not* correctable after T steps: this means there exists $x_a \neq x_b$, and error vectors $e_a^{(0)},\dots,e_a^{(T-1)}$ (supported on L_a with $|L_a| \leq q$) and $e_b^{(0)},\dots,e_b^{(T-1)}$ (supported on L_b, with $|L_b| \leq q$) such that $CA^t x_a + e_a^{(t)} = CA^t x_b + e_b^{(t)}$ for all $t \in \{0,\dots,T-1\}$. Now let $z = x_a - x_b \neq 0$. If we let $L = L_a \cup L_b$, then we have $|L| \leq 2q$, and we have for all $t \in \{0,\dots,T-1\}$, $\mathsf{supp}(CA^t z) \subset L$ which shows that (ii) does not hold. ∎

It is interesting to note the connection of the proposition above with the definition of a q-error-correcting code in the context of coding over the real numbers. A matrix $C \in \mathbb{R}^{p\times n}$ (with $p > n$) defines a q-error-correcting code (i.e., the code defined by C can correct q errors) if for any $z \neq 0$, $|\mathsf{supp}(Cz)| > 2q$ (see for example [14, §3]). This is precisely the condition we obtain from the previous proposition when $T = 1$ and there is no dynamics.

It is also interesting to observe that the proposition above shows that one cannot recover the initial state $x^{(0)}$ until the observability matrix given by

$$\begin{bmatrix} C \\ CA \\ \vdots \\ CA^{T-1} \end{bmatrix}$$

has rank n. Indeed, if the observability matrix has rank smaller

than n then it has a nontrivial kernel and there exists $z \neq 0$ such that $Cz = CAz = \cdots = CA^{T-1}z = 0$. This shows, by the above proposition, that "0 errors cannot be corrected", or in other words, that one cannot reconstruct $x^{(0)}$ even if there are no errors in the $y^{(0)}, \ldots, y^{(T-1)}$. The condition stated in proposition 2 can therefore be seen as a generalized condition for observability of a linear dynamical system when the observations are corrupted (as per the model considered here).

Observe also that the characterization of proposition 2 shows that the maximum number of correctable errors cannot increase beyond $T = n$ measurements. Indeed, this is a direct consequence of the Cayley-Hamilton theorem since we have for any z and for $t \geq n$, $\mathsf{supp}(CA^t z) \subset \mathsf{supp}(Cz) \cup \mathsf{supp}(CAz) \cup \cdots \cup \mathsf{supp}(CA^{n-1}z)$.

Finally, one can also directly see from the same proposition that the number of correctable errors is always less than $p/2$, for any T. In the next result we give a slightly more refined upper bound on the number of correctable errors as a function of T.

Proposition 3. Let $T \in \mathbb{N}\backslash\{0\}$ be such that $pT \geq n$, where $p \in \mathbb{N}\backslash\{0\}$ is the number of sensors in (1). If q errors are correctable after T steps, then necessarily $q < \frac{p - \lfloor (n-1)/T \rfloor}{2} \leq \frac{p - n/T + 1}{2}$.

Note that if $pT < n$ we are in the situation considered earlier where the observability matrix has rank smaller than n and where we cannot reconstruct $x^{(0)}$ even if there are no errors.

Proof: We show that there exists $z \neq 0$ such that $|\mathsf{supp}(Cz) \cup \mathsf{supp}(CAz) \cup \cdots \cup \mathsf{supp}(CA^{T-1}z)| \leq p - \lfloor (n-1)/T \rfloor$. Let L be any subset of $\{1, \ldots, p\}$ of cardinality $\lfloor (n-1)/T \rfloor$ (for example $L = \{1, \ldots, \lfloor (n-1)/T \rfloor\}$). Consider the linear operator $\Phi \colon z \in \mathbb{R}^n \mapsto (\mathcal{P}_L Cz, \mathcal{P}_L CAz, \ldots, \mathcal{P}_L CA^{T-1}z) \in \mathbb{R}^{|L|T}$, where $\mathcal{P}_L$ is the projection onto the components of L. The codomain of Φ is $\mathbb{R}^{|L|T}$, and since $|L| = \lfloor (n-1)/T \rfloor < n/T$, the codomain of Φ has dimension strictly less than n which means that Φ has a nontrivial kernel. This therefore shows that there exists a $z \neq 0$ such that $\mathsf{supp}(Cz) \cup \mathsf{supp}(CAz) \cup \cdots \cup \mathsf{supp}(CA^{T-1}z) \subseteq L^c$, and so $|\mathsf{supp}(Cz) \cup \mathsf{supp}(CAz) \cup \cdots \cup \mathsf{supp}(CA^{T-1}z)| \leq |L^c| = p - \lfloor (n-1)/T \rfloor$. ∎

We will now show that, when $T = n$, the upper bound given in the previous proposition is tight *generically*, that is, for "almost all" pairs of matrices (A, C).

Proposition 4. For almost all[1] pairs $(A, C) \in \mathbb{R}^{n \times n} \times \mathbb{R}^{p \times n}$ the number of correctable errors after $T = n$ steps is maximal and equal to $\lfloor (p-1)/2 \rfloor$.

Proof: For $i \in \{1, \ldots, p\}$, let $\mathcal{P}_{\{i\}} \colon \mathbb{R}^p \to \mathbb{R}$ be the projection map onto the i'th component ($\mathcal{P}_{\{i\}}$ is a $1 \times p$ matrix). Now let f_i be the map defined as:

$$f_i \colon (A, C) \in \mathbb{R}^{n \times n} \times \mathbb{R}^{p \times n} \mapsto \det \begin{bmatrix} \mathcal{P}_{\{i\}} C \\ \mathcal{P}_{\{i\}} CA \\ \vdots \\ \mathcal{P}_{\{i\}} CA^{n-1} \end{bmatrix} \in \mathbb{R}.$$

where the matrix in the argument of the determinant is an $n \times n$ matrix. Note that f_i is a polynomial in the entries of A and C that is not identically zero [to see this take for example $C = \mathcal{P}_{\{1,\ldots,p\}}$ to be the projection on the first p components and A to be a circular permutation matrix; then if z is such that $P_{\{i\}} Cz = 0, \ldots, P_{\{i\}} CA^{n-1}z = 0$ then necessarily $z = 0$, and so the $n \times n$ matrix in the definition of f_i has trivial kernel and so has nonzero determinant]. Hence the zero set Z_i of f_i, $Z_i = \{(A, C) \in \mathbb{R}^{n \times n} \times \mathbb{R}^{p \times n} \mid f_i(A, C) = 0\}$ has (Lebesgue) measure zero in $\mathbb{R}^{n \times n} \times \mathbb{R}^{p \times n}$ [15]. Hence $\cup_{i=1}^p Z_i$ has also measure zero. Now to conclude, note that for any $(A, C) \in (\cup_{i=1}^p Z_i)^c$, the number of correctable errors after $T = n$ is maximal. Indeed if $z \neq 0$, we have that for all i, $(P_{\{i\}} Cz, \ldots, P_{\{i\}} CA^{n-1}z) \neq 0_{\mathbb{R}^n}$ since $(A, C) \notin Z_i$, and so $|\mathsf{supp}(Cz) \cup \mathsf{supp}(CAz) \cup \cdots \cup \mathsf{supp}(CA^{n-1}z)|$ is maximal and equal to p. ∎

A. Computing the number of correctable errors

Note that even though we showed, that for *almost all* pairs (A, C) the maximum number of errors that can be corrected is maximal (equal to $\lfloor (p-1)/2 \rfloor$ for $T = n$), the problem of actually *computing* the number of errors that can be corrected for a given pair (A, C) after a given number of steps T is considered to be nontrivial. The simplest way to compute this number is in fact to consider every possible subset $K \subset \{1, \ldots, p\}$ and to check if the following $(p - |K|)T \times n$ matrix has a nontrivial kernel:

$$\begin{bmatrix} \mathcal{P}_{K^c} C \\ \mathcal{P}_{K^c} CA \\ \vdots \\ \mathcal{P}_{K^c} CA^{T-1} \end{bmatrix}.$$

If s is the cardinality of the smallest K for which this matrix has nontrivial kernel, then the maximum number of correctable errors is $\lfloor \frac{s-1}{2} \rfloor$ (indeed by definition of s we have that for any $z \neq 0$, $|\mathsf{supp}(Cz) \cup \cdots \cup \mathsf{supp}(CA^{T-1}z)| > s - 1$ and furthermore the inequality is tight since there exists $z \neq 0$ such that $|\mathsf{supp}(Cz) \cup \cdots \cup \mathsf{supp}(CA^{T-1}z)| = s$).

This algorithm is however combinatorial in nature and requires computing the rank of 2^p matrices in the worst-case. It is not clear if there exists a more efficient way to perform the computation [16][2].

[1] In other words, for all pairs (A, C) but a set of (Lebesgue) measure zero.

[2] In the special case $T = 1$ of error correction without dynamics, the number of errors that can be corrected is directly related to the *spark* of a matrix F that annihilates C, i.e., such that $FC = 0$ (see [13, §I.G]). The spark of a matrix F is the smallest number of columns that are linearly dependent. According to [16] it is still unknown whether computing spark is NP-hard.

IV. Decoding as an optimization problem

In this section we consider the problem of actually *constructing* a decoder that can correct the number of errors given by proposition 2.

Consider the decoder $D_0^T : (\mathbb{R}^p)^T \to \mathbb{R}^n$ defined such that $D_0^T(y^{(0)}, \dots, y^{(T-1)})$ is the optimal x solution of the following optimization problem:

$$\begin{aligned} \underset{x\in\mathbb{R}^n, K\subset\{1,\dots,p\}}{\text{minimize}} \quad & |K| \\ \text{subject to} \quad & \mathsf{supp}(y^{(t)} - CA^t x) \subset K \\ & \text{for } t \in \{0, \dots, T-1\}. \end{aligned} \tag{3}$$

If the optimization problem has more than one solution, we take $D_0^T(y^{(0)}, \dots, y^{(T-1)})$ to be any such solution. Observe that the decoder D_0^T looks for the smallest set K of attacked sensors that can explain the received data $y^{(0)}, \dots, y^{(T-1)}$. We show in the next proposition that the decoder D_0^T is, in some sense, optimal.

Proposition 5. Assume that q errors are correctable after T steps, i.e., that $|\mathsf{supp}(Cz) \cup \dots \cup \mathsf{supp}(CA^{T-1}z)| > 2q$ for all $z \in \mathbb{R}^n\backslash\{0\}$. Then the decoder D_0^T corrects q errors, i.e., for any $x^{(0)} \in \mathbb{R}^n$, and any $e^{(0)}, \dots, e^{(T-1)}$ in $\mathbb{R}^p$ such that $\mathsf{supp}(e^{(t)}) \subset K$ with $|K| \leq q$, we have $D_0^T(y^{(0)}, \dots, y^{(T-1)}) = x^{(0)}$ where $y^{(t)} = CA^t x^{(0)} + e^{(t)}$.

Proof: Let $x^{(0)}$ and the $e^{(t)}$'s be as in the proposition, with $\mathsf{supp}(e^{(t)}) \subset K$ and $y^{(t)} = CA^t + e^{(t)}$. Assume that the feasible point $(x^{(0)}, K)$ is not the optimal point for (3). Hence there exists $x_a \neq x^{(0)}$, and $e_a^{(0)}, \dots, e_a^{(T-1)}$ with $\mathsf{supp}(e_a^{(t)}) \subset K_a$ that generate the same sequence $y^{(0)}, \dots, y^{(T-1)}$ of observed values, with in addition, $|K_a| \leq |K| \leq q$. We therefore have two different initial conditions $x^{(0)} \neq x_a$ and two different error vectors corresponding to less than q attacked sensors that generate exactly the same sequence of observed values. This exactly means that q errors are *not* correctable after T steps which contradicts the assumption. ∎

The proposition above therefore shows that the decoder D_0^T is somehow the "best" decoder we can have since if any decoder can correct q errors, then D_0^T can as well. One issue however is that the optimization problem (3) is not practical since it is NP-hard in general. Indeed for the special case $T = 1$ (corresponding to the case of "static" error-correction over the reals mentioned earlier) the decoder becomes

$$\underset{x\in\mathbb{R}^n}{\text{minimize}} \quad \|y - Cx\|_{\ell_0} \tag{4}$$

(where $\|z\|_{\ell_0} = |\mathsf{supp}(z)|$) which is known to be NP-hard (see for example [14]).

However, in [13], Candes and Tao propose to replace the ℓ_0 "norm" by an ℓ_1 norm, thereby transforming the problem into a linear program that can be efficiently solved:

$$\underset{x\in\mathbb{R}^n}{\text{minimize}} \quad \|y - Cx\|_{\ell_1}.$$

It was then shown in [13] that if the matrix C satisfies certain conditions, then the solution of this linear program is the same as the one given by the ℓ_0 optimal decoder. In the next section we will operate this transformation in the context of our problem.

V. The ℓ_1 decoder

For $T \in \mathbb{N}\backslash\{0\}$, consider the linear map $\Phi^{(T)}$ defined by:

$$\begin{aligned} \Phi^{(T)} \colon \mathbb{R}^n &\to \mathbb{R}^{p\times T} \\ x &\mapsto \begin{bmatrix} Cx & | & CAx & | & \dots & | & CA^{T-1}x \end{bmatrix}. \end{aligned}$$

Furthermore, if $y^{(0)}, \dots, y^{(T-1)} \in \mathbb{R}^p$, let $Y^{(T)}$ the $p \times T$ matrix formed by catenening the $y^{(t)}$'s in columns:

$$Y^{(T)} = \begin{bmatrix} y^{(0)} & | & y^{(1)} & | & \dots & | & y^{(T-1)} \end{bmatrix} \in \mathbb{R}^{p\times T}.$$

Recall that for a matrix $M \in \mathbb{R}^{p\times T}$ with rows $M_1, \dots, M_p \in \mathbb{R}^T$ the ℓ_0 "norm" of M is the number of nonzero rows in M:

$$\|M\|_{\ell_0} = |\mathsf{rowsupp}(M)| = |\{i \in \{1, \dots, p\} \mid M_i \neq 0_{\mathbb{R}^T}\}|.$$

Observe that the optimal decoder D_0^T introduced in the previous section can be written as:

$$D_0^T(y^{(0)}, \dots, y^{(T-1)}) = \underset{x\in\mathbb{R}^n}{\operatorname{argmin}} \|Y^{(T)} - \Phi^{(T)}x\|_{\ell_0}.$$

As we saw in the previous section, this decoder finds the minimum number of attacked sensors that can explain the received data $y^{(0)}, \dots, y^{(T-1)}$.

Analogously to [13], we can define an ℓ_1 decoder in which, instead of minimizing the number of nonzero rows, we will minimize the sum of the magnitudes of each row. More specifically, if we measure the magnitude of a row by its ℓ_r norm in $\mathbb{R}^T$ (for $r \geq 1$), then we obtain the following decoder $D_{1,r}^T$:

$$D_{1,r}^T(y^{(0)}, \dots, y^{(T-1)}) = \underset{x\in\mathbb{R}^n}{\operatorname{argmin}} \|Y^{(T)} - \Phi^{(T)}x\|_{\ell_1/\ell_r} \tag{5}$$

where, by definition, $\|M\|_{\ell_1/\ell_r}$ is the sum of the ℓ_r norms of the rows of the matrix M:

$$\|M\|_{\ell_1/\ell_r} = \sum_{i=1}^{p} \|M_i\|_{\ell_r}.$$

Note that the optimization problem in (5) is convex and can be efficiently solved. Also note that such "mixed" ℓ_1/ℓ_r norms were also used in the compressed sensing literature in the context of joint-sparse and block-sparse signal recovery [17].

A. Number of errors correctable by the ℓ_1/ℓ_r decoder

We saw in proposition 2 that the number of errors that can be corrected by the optimal ℓ_0 decoder D_0^T is equal to the largest number q such that $|\mathsf{supp}(Cz) \cup \mathsf{supp}(CAz) \cup \dots \cup \mathsf{supp}(CA^{T-1}z)| > 2q$ for all $z \neq 0$.

The next proposition now characterizes the maximum number of errors that can be corrected by the ℓ_1/ℓ_r decoder $D_{1,r}^T$.

Proposition 6. The following are equivalent:
(i) The decoder $D_{1,r}^T$ can correct q errors after T steps.
(ii) For all $K \subset \{1, \dots, p\}$ with $|K| = q$ and for all $G = \Phi^{(T)}z$ with $z \in \mathbb{R}^n\backslash\{0\}$ we have:

$$\sum_{i\in K} \|G_i\|_{\ell_r} < \sum_{i\in K^c} \|G_i\|_{\ell_r}. \tag{6}$$

Proof: **(i)** $\Rightarrow$ **(ii)**: Suppose for the sake of contradiction that (ii) does not hold. Then there exists $K \subset \{1,\dots,p\}$ with $|K| = q$, and $G = \Phi^{(T)}z \in \mathbb{R}^{p\times T}$ with $z \neq 0$ such that $\sum_{i\in K}\|G_i\|_{\ell_r} \geq \sum_{i\in K^c}\|G_i\|_{\ell_r}$. Let $x^0 = 0$ and define the K-supported error vectors $e^{(t)}$, for $t \in \{0,\dots,T-1\}$ by:

$$e_i^{(t)} = \begin{cases} G_{i,t} & \text{if } i \in K \\ 0 & \text{otherwise} \end{cases}$$

Now consider $y^{(t)} = CA^t x^0 + e^{(t)} = e^{(t)}$ and let $Y^{(T)}$ be as before the $p \times T$ matrix obtained by concatenating the $y^{(t)}$'s in columns. Note that $\mathsf{rowsupp}(Y^{(T)}) = K$, and that $Y_i^{(T)} = (\Phi^{(T)}z)_i$ for all $i \in K$. We will now show that the objective function for (5) at $z \neq 0$ is smaller than at $x^0 = 0$, which will show that the decoder $D_{1,r}^T$ fails to reconstruct $x^{(0)}$ from the $y^{(t)}$'s. This will show that (i) is not true. Indeed we have:

$$\begin{aligned}\|Y^{(T)} - \Phi^{(T)}z\|_{\ell_1/\ell_r} &= \sum_{i=1}^n \|(Y^{(T)} - \Phi^{(T)}z)_i\|_{\ell_r} \\ &= \sum_{i\in K^c}\|G_i\|_{\ell_r} \leq \sum_{i\in K}\|G_i\|_{\ell_r} \\ &= \sum_{i=1}^n \|(Y^{(T)} - \Phi^{(T)}x^0)_i\|_{\ell_r} \\ &= \|Y^{(T)} - \Phi^{(T)}x^0\|_{\ell_1/\ell_r}.\end{aligned}$$

(ii) $\Rightarrow$ **(i)**: We again resort to contradiction. Suppose that (i) is not true. This means there exists $x^{(0)}$, and $e^{(0)},\dots,e^{(T-1)}$ with $\mathsf{supp}(e^{(t)}) \subset K$ with $|K| \leq q$ such that $D_{1,r}^T(y^{(0)},\dots,y^{(T-1)}) \neq x^{(0)}$ where $y^{(t)} = CA^t x^{(0)} + e^{(t)}$ (i.e., the decoder $D_{1,r}^T$ fails to reconstruct $x^{(0)}$ from the $y(t)$'s). By definition of the decoder $D_{1,r}^T$, this means that there exists $\tilde{x} \neq x^{(0)}$ that achieves a smaller ℓ_1/ℓ_r objective than $x^{(0)}$:

$$\sum_{i=1}^n \|(Y^{(T)} - \Phi^{(T)}\tilde{x})_i\|_{\ell_r} \leq \sum_{i=1}^n \|(Y^{(T)} - \Phi^{(T)}x^{(0)})_i\|_{\ell_r}.$$

Now let $z = \tilde{x} - x^{(0)} \neq 0$, and let $G = \Phi^{(T)}z = U - V$ with $U = Y^{(T)} - \Phi^{(T)}x^{(0)}$ and $V = Y^{(T)} - \Phi^{(T)}\tilde{x}$. We have

$$\sum_{i\in K}\|G_i\|_{\ell_r} = \sum_{i\in K}\|U_i - V_i\|_{\ell_r} \geq \sum_{i\in K}\|U_i\|_{\ell_r} - \|V_i\|_{\ell_r}$$

Now since $\mathsf{rowsupp}(U) \subset K$, and since $\tilde{x}$ achieves a smaller ℓ_1/ℓ_r objective than x^0, we have $\sum_{i\in K}\|U_i\|_{\ell_r} = \sum_{i=1}^n \|U_i\|_{\ell_r} \geq \sum_{i=1}^n \|V_i\|_{\ell_r}$. Hence we have

$$\begin{aligned}\sum_{i\in K}\|G_i\|_{\ell_r} &\geq \sum_{i=1}^n \|V_i\|_{\ell_r} - \sum_{i\in K}\|V_i\|_{\ell_r} \\ &= \sum_{i\in K^c}\|V_i\|_{\ell_r} = \sum_{i\in K^c}\|G_i\|_{\ell_r}\end{aligned}$$

where the last equality is because $\mathsf{rowsupp}(U) \subset K$. Hence (ii) is not true. ∎

Observe that, as expected, if the ℓ_1/ℓ_r decoder can correct q errors, then the ℓ_0 decoder can correct q errors as well. Indeed, if we assume the opposite, then by proposition 2 there exists $z \neq 0$ such that $|\mathsf{supp}(Cz) \cup \cdots \cup \mathsf{supp}(CA^{T-1}z)| \leq 2q$, which is equivalent to saying that $|\mathsf{rowsupp}(\Phi^{(T)}z)| \leq 2q$. Now let $G = \Phi^{(T)}z$ and let K be the q rows of G with the largest ℓ_r norms, then we clearly have $\sum_{i\in K}\|G_i\|_{\ell_r} \geq \sum_{i\in K^c}\|G_i\|_{\ell_r}$, which contradicts the condition of the previous proposition.

As a matter of fact, the condition of the previous proposition (for the ℓ_1/ℓ_r decoder) is in some sense a more quantitative version of the condition of proposition 2 for the ℓ_0 decoder. The two conditions guarantee that the row components of $\Phi^{(T)}z$ are sufficiently spread and are not too concentrated on a small subset of the rows.

As an illustration, consider the simple example where the number of sensors is $p = n$ and $C = \mathrm{Id}_{\mathbb{R}^n}$ (i.e., we have one sensor per component of the state $x \in \mathbb{R}^n$) and assume that A is a cyclic permutation, say:

$$A = \begin{bmatrix} 0 & 1 & \dots & 0 & 0 \\ 0 & 0 & 1 & \dots & 0 \\ \vdots & \vdots & & \ddots & 0 \\ 0 & 0 & \dots & 0 & 1 \\ 1 & 0 & 0 & \dots & 0 \end{bmatrix}$$

It is easy to see that after $T = n$, the rows of the matrix $\Phi^{(n)}z = \begin{bmatrix} z & Az & \dots & A^{n-1}z \end{bmatrix}$ are identical up to a permutation, and so the ℓ_r norm of any two rows of $\Phi^{(T)}z$ are equal. This shows that for any subset K of rows with $|K| < n/2$, we have $\sum_{i\in K}\|(\Phi^{(n)}z)_i\|_{\ell_r} < \sum_{i\in K^c}\|(\Phi^{(n)}z)_i\|_{\ell_r}$, which shows that the ℓ_1/ℓ_r decoder can correct a maximal number of errors after n steps, namely, $\lfloor (n-1)/2 \rfloor$.

Finally note that the condition of proposition 6 for the ℓ_1/ℓ_r decoder corresponds to the well-known "nullspace property" in compressed sensing and sparse signal recovery [18].

B. A more tractable sufficient condition

The number of errors that can be corrected by the ℓ_1/ℓ_r decoder is the largest q that satisfies the condition of proposition 6. The proposed condition however does not directly give a way to *compute* this number since it involves checking an inequality for *every* $z \in \mathbb{R}^n\backslash\{0\}$. In this section we will propose one way to obtain simpler sufficient conditions for the ℓ_1/ℓ_r decoder to correct q errors.

First observe that for a given $K \subset \{1,\dots,p\}$ and a given $z \in \mathbb{R}^n\backslash\{0\}$, condition (6) of proposition 6 can be rewritten as follows:

$$\begin{aligned}(6) &\Leftrightarrow \sum_{i\in K}\|(\Phi^{(T)}z)_i\|_{\ell_r} < \sum_{i\in K^c}\|(\Phi^{(T)}z)_i\|_{\ell_r} \\ &\Leftrightarrow \frac{\|(\Phi^{(T)}z)_K\|_{\ell_1/\ell_r}}{\|(\Phi^{(T)}z)_{K^c}\|_{\ell_1/\ell_r}} < 1,\end{aligned}$$

where we have used the notation $(\Phi^{(T)}z)_K \in \mathbb{R}^{|K|\times T}$ for the $|K| \times T$ matrix obtained from $\Phi^{(T)}z \in \mathbb{R}^{p\times T}$ by keeping only the rows in K (similarly for $(\Phi^{(T)}z)_{K^c}$). Using this notation, we can say that the ℓ_1/ℓ_r decoder can correct q errors if and only if the following condition holds:

$$\sup_{\substack{K\subset\{1,\dots,p\} \\ |K|=q}} \sup_{z\in\mathbb{R}^n\backslash\{0\}} \frac{\|(\Phi^{(T)}z)_K\|_{\ell_1/\ell_r}}{\|(\Phi^{(T)}z)_{K^c}\|_{\ell_1/\ell_r}} < 1 \qquad (7)$$

In order to simplify this condition, we will look for a way to upper bound the expression $\frac{\|(\Phi^{(T)}z)_K\|_{\ell_1/\ell_r}}{\|(\Phi^{(T)}z)_{K^c}\|_{\ell_1/\ell_r}}$ uniformly in z so that the above inequality becomes more computationally accessible. When the chosen norm is $\ell_r = \ell_2$, a very simple way is to use the extreme singular values of suitably chosen operators which allow us to write $\sigma_{\min}\|z\|_2 \leq \|Lz\|_2 \leq \sigma_{\max}\|z\|_2$ when L is a linear map and $\sigma_{\min}$ and $\sigma_{\max}$ are the smallest and largest singular values of L. Of course this will give us conditions that are not tight in general, but the advantage is that they will be more computationally tractable.

Recall that, by definition of the ℓ_1/ℓ_r norm, we have $\|(\Phi^{(T)}z)_K\|_{\ell_1/\ell_r} = \sum_{i\in K}\|(\Phi^{(T)}z)_i\|_{\ell_r}$. Let $\Phi_i^{(T)}$ be the linear map from $\mathbb{R}^n$ to $\mathbb{R}^T$ such that $\Phi_i^{(T)}z = (\Phi^{(T)}z)_i$ for all $z \in \mathbb{R}^n$. A matrix representation of this map is given by:

$$\begin{bmatrix} \mathcal{P}_{\{i\}}C \\ \mathcal{P}_{\{i\}}CA \\ \vdots \\ \mathcal{P}_{\{i\}}CA^{T-1} \end{bmatrix} \in \mathbb{R}^{T\times n},$$

where, as before, $\mathcal{P}_{\{i\}}$ is the projection map onto the i'th component. Now let $K \subset \{1,\dots,p\}$ such that $|K| = q$ be fixed. To bound the numerator of the expression $\frac{\|(\Phi^{(T)}z)_K\|_{\ell_1/\ell_r}}{\|(\Phi^{(T)}z)_{K^c}\|_{\ell_1/\ell_r}}$ we can use the ℓ_r-operator norms $\|\Phi_i^{(T)}\|_{\ell_r}$ of the linear maps $\Phi_i^{(T)}$ and we have:

$$\|(\Phi^{(T)}z)_K\|_{\ell_1/\ell_r} = \sum_{i\in K}\|\Phi_i^{(T)}z\|_{\ell_r} \leq \sum_{i\in K}\|\Phi_i^{(T)}\|_{\ell_r}\|z\|_{\ell_r}$$

If we call $\beta = \max_{i=1,\dots,p}\|\Phi_i^{(T)}\|_{\ell_r}$ then we can write

$$\|(\Phi^{(T)}z)_K\|_{\ell_1/\ell_r} \leq q\beta\|z\|_{\ell_r}, \tag{8}$$

since $|K| = q$. Note that when $\ell_r = \ell_2$, $\|\Phi_i^{(T)}\|_{\ell_r}$ is the largest singular value of $\Phi_i^{(T)}$. We now turn to the problem of finding a lower bound for the denominator of $\frac{\|(\Phi^{(T)}z)_K\|_{\ell_1/\ell_r}}{\|(\Phi^{(T)}z)_{K^c}\|_{\ell_1/\ell_r}}$. For this, we will directly consider the special case $\ell_r = \ell_2$, since we can then write $\|\Phi_i^{(T)}z\|_{\ell_2} \geq \sigma_{\min}(\Phi_i^{(T)})\|z\|_{\ell_2}$, where $\sigma_{\min}(\Phi_i^{(T)})$ is the smallest singular value of the linear map $\Phi_i^{(T)}$. Now if we call $\alpha = \min_{i=1,\dots,p}\sigma_{\min}(\Phi_i^{(T)})$, we get that

$$\|(\Phi^{(T)}z)_{K^c}\|_{\ell_1/\ell_2} = \sum_{i\in K^c}\|\Phi_i^{(T)}z\|_{\ell_2} \geq (p-q)\alpha\|z\|_{\ell_2}. \tag{9}$$

Hence if we combine (8) and (9), we get:

$$\sup_{z\in\mathbb{R}^n\setminus\{0\}}\frac{\|(\Phi^{(T)}z)_K\|_{\ell_1/\ell_2}}{\|(\Phi^{(T)}z)_{K^c}\|_{\ell_1/\ell_2}} < \frac{q\beta}{(p-q)\alpha}. \tag{10}$$

Note that the right-hand side of this inequality does not depend on the set K. It follows that a sufficient condition for q errors to be correctable by the ℓ_1/ℓ_2 decoder is that $q\beta/((p-q)\alpha) < 1$, i.e., $q < \frac{p\alpha}{\alpha+\beta}$. In other words, the ℓ_1/ℓ_2 can correct at least $\lceil\frac{p\alpha}{\alpha+\beta} - 1\rceil$ errors.

Observe that for the simple example mentioned in the previous section where $C = \mathrm{Id}_{\mathbb{R}^n}$ and A is a circular permutation matrix we obtain that after $T = n$ steps, $\lceil n/2 - 1\rceil$ errors are correctable since in this case have $\alpha = \beta = 1$. For this particular choice of matrices (A, C) the result is therefore tight.

It is clear however from the derivation (and as we mentioned in the beginning) that this bound is in general quite loose. Observe for example that this result is only useful when $T \geq n$. Indeed, if $T < n$, then for $i \in \{1,\dots,p\}$ the linear map $\Phi_i^{(T)}\colon \mathbb{R}^n \to \mathbb{R}^T$ has a nontrivial kernel and so $\sigma_{\min}(\Phi_i^{(T)}) = 0$ which means that $\alpha = 0$. We are currently working on improving this bound.

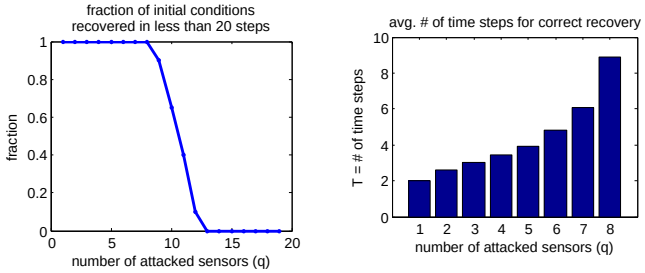

(a) (b)

Fig. 1. (a) Fraction of initial conditions that were correctly recovered in less than $T = 20$ time steps, for different values of q. (b) Average number of time steps it took to correctly recover the initial state, as a function of the number of corrupted components.

VI. Numerical simulations

In this section we show the performance of the proposed decoding algorithm first on a random toy example and then on a more realistic system modeling an electric power network.

A. Random system

We first consider the ℓ_1/ℓ_2 decoder on a system of size $n = 30$, $p = 20$ where $A \in \mathbb{R}^{30\times 30}$ and $C \in \mathbb{R}^{20\times 30}$ have iid Gaussian entries. For different values q of attacked sensors, we tested the decoder on 20 random initial conditions and randomly chosen sets of attacked sensors $K \subset \{1,\dots,p\}$ with $|K| = q$: Figure 1a shows the fraction of initial conditions that were correctly recovered in less than $T = 20$ time steps for the different values of q. We see that for q less than 8 all the initial conditions were correctly recovered in less than $T = 20$ time steps. Figure 1b shows the number of time steps that it took in average to correctly recover the initial state, as a function of the number of corrupted components q.

For each simulation, the error values (i.e., the values injected by the attacker in the components K) were chosen randomly, and their magnitudes were five times larger than the magnitude of the state. The matrix A was appropriately scaled so it has a spectral radius of 1. The optimization problems were solved using `CVX` [19].

B. Electric power network

In this section we will apply the proposed decoding algorithm on a model of an electric power network and more specifically on the IEEE 14-bus power network [20]. The network, depicted in figure 2 is composed of 5 synchronous

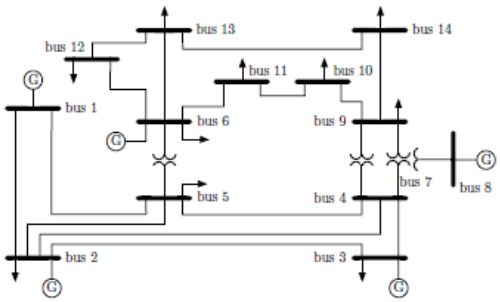

Fig. 2. IEEE 14-bus power network [9].

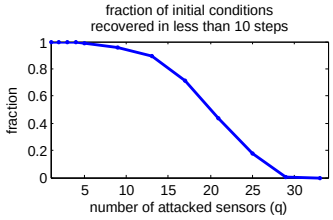

Fig. 3. Fraction of initial conditions that were correctly recovered in less than $T = 10$ steps for the IEEE 14-bus example. For each value of q, 200 simulations were carried out with different initial conditions and different sets of attacked sensors.

generators and a total of 14 buses. The system is represented by $2 \times 5 = 10$ states giving the rotor angles δ_i and the frequencies $\omega_i = d\delta_i/dt$ of each generator. Under some simplifying assumptions the evolution of the system can be captured by a linear difference equation corresponding to the linearized swing equations (see [21] for the derivation of the equations). We assume, like in [9], that $p = 35$ sensors are deployed and measure at every time step the real power injections at every bus (14 sensors), the real power flows along every branch (20 sensors), and the rotor angle at generator 1 (1 sensor).

Following [9], we assume that all, but the sensor measuring the rotor angle can be attacked by a malicious agent. For different values of q between 1 and 34, we ran 200 simulations with different sets of attacked sensors K of cardinality q, and different initial conditions $x^{(0)}$. Figure 3 shows the number of simulations (out of the 200) where the state $x^{(0)}$ was correctly recovered using the ℓ_1/ℓ_∞ decoder in less than $T = 10$ steps. Observe that for $q \leq 4$ the success rate of the decoder was 100%. Furthermore when $q \leq 12$ the decoder correctly recovers the state in more than 90% of the cases. These simulations show that the ℓ_1/ℓ_r decoder works very well in practice.

VII. Discussion

The problem of computing the number of errors tolerated by the ℓ_1/ℓ_r decoder, using the characterization of proposition 6, is considered to be nontrivial. Some researchers in the compressed sensing community, for example in [22], have proposed approximate algorithms to verify the "nullspace property" which, as we mentioned earlier, is directly connected to proposition 6. It would be interesting to know if these algorithms can be adapted to the problem considered here.

Another important line of research is to find efficient ways to solve the optimization problem of the ℓ_1/ℓ_r decoder. Indeed, the problem size grows with time T and can become prohibitively large for large systems with real-time constraints.

Finally, it would be interesting to study the *robustness* of the proposed decoder against noise in the unattacked sensors and disturbances in the state-evolution equation.

Acknowledgment

We would like to thank Fabio Pasqualetti for the data used in the IEEE 14-bus power network example.

References

[1] A. Cardenas, S. Amin, B. Sinopoli, A. Giani, A. Perrig, and S. S. Sastry, "Challenges for securing cyber physical systems," in *Workshop on Future Directions in Cyber-physical Systems Security*. DHS, July 2009. [Online]. Available: http://chess.eecs.berkeley.edu/pubs/601.html

[2] "Critical infrastructure protection: Challenges and efforts to secure control systems," United States General Accounting Office (GAO), report to the Congress, GAO-04-354, March 2004.

[3] N. Adam, "Workshop on future directions in cyber-physical systems security," Report on workshop organized by Department of Homeland Security (DHS), January 2010.

[4] M. Blanke, M. Kinnaert, J. Lunze, and M. Staroswiecki, *Diagnosis and Fault-tolerant Control*. Springer, 2006.

[5] K. Zhou and J. Doyle, *Essentials of robust control*. Prentice Hall New Jersey, 1998, vol. 104.

[6] T. Başar and G. Olsder, *Dynamic noncooperative game theory*. SIAM, 1999.

[7] L. Schenato, B. Sinopoli, M. Franceschetti, K. Poolla, and S. Sastry, "Foundations of control and estimation over lossy networks," *Proceedings of the IEEE*, vol. 95, no. 1, pp. 163–187, 2007.

[8] A. Gupta, C. Langbort, and T. Basar, "Optimal control in the presence of an intelligent jammer with limited actions," in *IEEE CDC 2010*.

[9] F. Pasqualetti, F. Dörfler, and F. Bullo, "Cyber-physical attacks in power networks: Models, fundamental limitations and monitor design," *arXiv:1103.2795*, 2011.

[10] S. Sundaram and C. Hadjicostis, "Distributed function calculation via linear iterative strategies in the presence of malicious agents," *IEEE Trans. Autom. Control*, July 2011, to appear.

[11] S. Sundaram, M. Pajic, C. Hadjicostis, R. Mangharam, and G. Pappas, "The wireless control network: monitoring for malicious behavior," in *IEEE CDC 2010*.

[12] F. Pasqualetti, A. Bicchi, and F. Bullo, "Consensus computation in unreliable networks: A system theoretic approach," *arXiv:1007.2738*, 2010.

[13] E. Candes and T. Tao, "Decoding by linear programming," *IEEE Trans. Inf. Theory*, vol. 51, no. 12, pp. 4203–4215, 2005.

[14] V. Guruswami, J. Lee, and A. Wigderson, "Euclidean sections of ℓ_1^n with sublinear randomness and error-correction over the reals," in *Proceedings of RANDOM 2008*.

[15] J. Dion, C. Commault, and J. van der Woude, "Generic properties and control of linear structured systems: a survey," *Automatica*, 2003.

[16] L. Lim and P. Comon, "Multiarray signal processing: Tensor decomposition meets compressed sensing," *Comptes Rendus Mecanique*, 2010.

[17] Y. Eldar and H. Bolcskei, "Block-sparsity: Coherence and efficient recovery," in *IEEE ICASSP 2009*.

[18] M. Davenport, M. Duarte, Y. Eldar, and G. Kutyniok, "Introduction to compressed sensing," in *Compressed Sensing: Theory and Applications*. Cambridge University Press, 2011, pp. 1–68.

[19] M. Grant and S. Boyd, "CVX: Matlab software for disciplined convex programming, version 1.21," http://cvxr.com/cvx, Apr. 2011.

[20] R. Christie, "Power systems test case archive," http://www.ee.washington.edu/research/pstca/, 2000.

[21] F. Pasqualetti, A. Bicchi, and F. Bullo, "A graph-theoretical characterization of power network vulnerabilities," in *IEEE ACC 2011*.

[22] A. d'Aspremont and L. El Ghaoui, "Testing the nullspace property using semidefinite programming," *Mathematical Programming*, 2008.