THEORY OF COMPUTING, Volume 16 (986), 2020, pp. 1–23
www.theoryofcomputing.org

Distributed Corruption Detection in Networks

Noga Alon* Elchanan Mossel† Robin Pemantle‡

Received November 29, 2015; Revised April 18, 2019, March 9, 2020; Published March 31, 2020

Abstract: We consider the problem of distributed corruption detection in networks. In this model each node of a directed graph is either truthful or corrupt. Each node reports the type (truthful or corrupt) of each of its outneighbors. If it is truthful, it reports the truth, whereas if it is corrupt, it reports adversarially. This model, first considered by Preparata, Metze and Chien in 1967, motivated by the desire to identify the faulty components of a digital system by having the other components checking them, became known as the PMC model. The main known results for this model characterize networks in which *all* corrupt (that is, faulty) nodes can be identified, when there is a known upper bound on their number. We are interested in networks in which a *large fraction* of the nodes can be classified. It is known that in the PMC model, in order to identify all corrupt nodes when their number is t, all in-degrees have to be at least t. In contrast, we show that in d regular-graphs with strong expansion properties, a $1-O(1/d)$ fraction of the corrupt nodes, and a $1-O(1/d)$ fraction of the truthful nodes can be identified, whenever there is a majority of truthful nodes. We also observe that if the graph is very far from being a good expander, namely, if the deletion

*Research supported in part by NSF grant DMS-1855464, ISF grant 281/17, BSF grant 2018267 and the Simons Foundation.
†Research supported in part by NSF awards DMS-1737944, ONR N00014-16-1-2227, N00014-17-1-2598, ARO MURI W911NF1910217 and Simons Investigator award (622132).
‡Research supported in part by NSF grant # DMS-1209117.

ACM Classification: C.2.m, G.2.2

AMS Classification: 05C90, 68R10, 94C12

Key words and phrases: distributed computing, expander graphs, PMC model, graph separators, network testing

DOI: 10.4086/toc.2020.v016a986

of a small set of nodes splits the graph into small components, then no corruption detection is possible even if most of the nodes are truthful. Finally we discuss the algorithmic aspects and the computational hardness of the problem.

1 Introduction

We study the problem of *corruption detection* in networks. Given a network of agents, a subset of whom are corrupt and the rest are truthful, our goal is to find as many truthful and corrupt agents as possible. Neighboring agents audit each other. We assume that truthful agents report the type of their neighbors accurately. We make no assumption on the report of corrupt agents. For example, two corrupt neighbors may collude and report each other as non-corrupt. Similarly, a corrupt agent may prefer to report the status of some of its neighbors accurately, hoping that this will establish a truthful record for itself. We permit that the corrupt agents coordinate their actions in an arbitrary fashion.

The corruption model studied here is identical to the model of diagnosable systems that was introduced by Perparata, Metze and Chien [23] as a model of a digital system with many components that can fail. It is assumed that components can test some other components. The goal in [23] and in follow-up work, including [12, 15, 16], is to characterize networks that can detect all corrupt nodes as long as their number does not exceed a given value. Similar models have been introduced and studied in other areas of computer science, including Byzantine computing [17] and intrusion detection in the security community [21]. See also the survey [26]. A byproduct of this line of research is the VLSI chips puzzle discussed in [8, Problems 4-6]. In this puzzle there are n supposedly identical VLSI chips that in principle are capable of testing each other. A basic test involves two chips, each chip tests the other and reports whether it is good or bad. A good chip always gives an accurate report, but the answer of a bad chip cannot be trusted. The objective is to find a good chip, or all good chips, assuming more than half of the chips are good, using the minimum possible number of tests. This is (an adaptive version of) the corruption detection problem on a complete graph.

The original motivation for our work has been corruption detection in social and economic networks, where the main objective is to understand the structure of networks that enable one to identify as many of the corrupt nodes and as many of the truthful ones as possible. Our goal is to understand theoretically which network structures are more amenable to corruption and which are more robust against it. In particular, is it possible to design sparse networks that allow to detect corruption even if there are many corrupt nodes? Social scientists have studied many aspects of corruption in networks, see, e. g., [22, 25, 10]. However, to the best of our knowledge, prior to this work, the effect of the network structure on corruption detection has not been systematically studied.

As pointed out in follow-up work [14], while "corruption *is* prevalent in many real-world networks, [...] in many scenarios it is not easy to pinpoint even a single truthful node" in spite of our positive theoretical results. "One reason for that is that some of the assumptions do not seem to hold in some real-world networks. For example, we assume that audits from the truthful nodes are not only non-malicious, but also perfectly reliable. In practice this assumption is unlikely to be true: many truthful nodes could be non-malicious but simply unable to audit their neighbors accurately. Further assumptions that may not hold in some scenarios include the notion of a central agency that is both uncorrupted and has access to reports from every agent and possibly even the assumption that the number of corrupt nodes is less

than $|V|/2$." In addition there are many constraints on the network that may prohibit it to be an expander. Despite these shortcomings of our model, our work points to ideal conditions that allow corruption detection. While it is perhaps unreasonable to assume that one imposes an expander auditing structure among government agencies, imposing such a structure among more equal entities may be more realistic. Furthermore, our results for directed graphs suggest such structures even in cases where the auditing relation is not symmetric.

1.1 Formal definitions and main results

Definition 1.1 (Digraph, Oriented Graph, Graph)**.** For a set V, let $V^{(2)}$ denote the set of ordered pairs of distinct elements. A *digraph* $G = (V,E)$ consists of a set V of nodes ("agents") and a set $E \subset V^{(2)}$ of directed edges. If E contains no anti-parallel directed edges (a pair of edges (u,v) and (v,u)), then G is an *oriented graph*. If E contains a directed edge (u,v) if and only if it contains (v,u), G is an *undirected graph* or simply a *graph*.

Definition 1.2 (Type, Report)**.** Consider a digraph $G = (V,E)$ along with a function $\tau : V \to \{\mathtt{c},\mathtt{t}\}$ that assigns each node a *type*. We call τ a *truthfulness assignment* to G. We call a node v *truthful* if $\tau(v) = \mathtt{t}$ and corrupt if $\tau(v) = \mathtt{c}$. We write $T = \tau^{-1}(\mathtt{t})$ for the set of truthful nodes and $B = \tau^{-1}(\mathtt{c})$ for the set of corrupt nodes, so that $V = T \sqcup B$ is a partition of the set of nodes. A *report* is a function $\rho : E \to \{\mathtt{c},\mathtt{t}\}$. A report $\rho : E \to \{\mathtt{c},\mathtt{t}\}$ is *compatible* with the truthfulness assignment τ if for each truthful node $u \in T$ and each directed edge $(u,v) \in E$ we have $\rho(u,v) = \tau(v)$. We call $\rho(u,v)$ the type of v reported by u. We call τ a *valid* truthfulness assignment if $|T| > |B|$. We say that a report ρ is *feasible* if it is compatible with at least one valid truthfulness assignment.

Since it is common to use the letter C for constants, we prefer to denote the set of corrupt nodes by B (the set of "bad" nodes).

The question we address is, under what conditions on the digraph G and the number of truthful nodes it is possible to determine the truthfulness status of most nodes with certainty. It is easy to see that this is impossible if $|T| \leq |B|$. Indeed, if $V = V_1 \cup V_2 \cup W$ is a partition of V into 3 pairwise disjoint sets where $|V_1| = |V_2|$ (and W may be empty), then the corrupt nodes can ensure that all the reports in the two scenarios $T = V_1$, $B = V_2 \cup W$ and $T = V_2$, $B = V_1 \cup W$ will be identical. As there is no common truthful node in these two possibilities, no algorithm can locate a truthful node with no error.

Our main result is that if the graph is a good bounded-degree directed expander, in the sense described below, and we have a majority of truthful nodes, then it is possible to determine the truthfulness status of most of the nodes.

We recall the definition of regular spectral expander graphs.

Definition 1.3. A graph G is an (n,d,λ)*-graph* if it is d-regular, has n nodes and all its eigenvalues besides the top one are in absolute value at most λ. Such a graph with $\lambda = 2\sqrt{d-1}$ is called a *Ramanujan graph*.

Explicit constructions of Ramanujan graphs were given by Lubotzky, Phillips and Sarnak [19] and by Margulis [20]. For the known connection between the expansion and pseudo-random properties of graphs and their eigenvalues, see, e. g., [3], [13] and the references therein.

First we consider the somewhat simpler case of *undirected graphs*.

The main result for the undirected case is the following. A more general version is stated as Theorem 2.1 in subsection 2.1.

Theorem 1.4. *Let $G=(V,E)$ be a Ramanujan graph, that is, an (n,d,λ) graph with $\lambda = 2\sqrt{d-1}$, and suppose $d \geq 100$. Then, for every report ρ, there exist sets T' and B' of nodes such that for every valid truthfulness assignment τ compatible with ρ we have*

$$T' \subset T, \quad B' \subset B, \quad |T \setminus T'| \leq \frac{32}{d}|B|, \quad |B \setminus B'| \leq \frac{32}{d}|B|, \tag{1.1}$$

where $T = \tau^{-1}(\mathsf{t})$ and $B = \tau^{-1}(\mathsf{c})$. Moreover, there is a linear-time algorithm that finds subsets T' and B' satisfying (1.1) for every truthfulness assignment τ compatible with ρ which satisfies $|T| > (1+12/d)(n/2)$.

Recall that $|T| > |B|$ in the theorem above since τ is valid. We make a few remarks on the conclusions of the theorem.

1. Given any graph G with all degrees bounded by d, with at least $2d+1$ corrupt nodes, then there is a set of $\lfloor |B|/(2d+1) \rfloor$ corrupt nodes and a set of $\lfloor |B|/(2d+1) \rfloor$ truthful nodes that cannot be classified even if in addition to the report ρ, we are given the number of corrupt nodes. Thus, the fraction of each type of nodes we fail to classify is tight up to a constant factor. To see that this is the case, consider the following selection of b corrupt nodes. Choose an arbitrary node v_1 in the graph and set all its neighbors $N(v_1)$ to be corrupt, then choose $v_2 \notin \{v_1\} \cup N(v_1)$ and set all of its neighbors to be corrupt, then choose one of v_1, v_2 to be corrupt and the other to be truthful. Continue in a similar fashion, defining $N(v_3), N(v_4)$ to be corrupt and choosing one of v_3, v_4 to be corrupt and the other truthful. Continue until the number of corrupt nodes left, b', satisfies $b' < 2d+1$. Set b' of the remaining nodes to be corrupt. Let $r(v,u) = \mathsf{c}$, whenever v is corrupt. It is then clear that it is impossible to decide which of the v_i is corrupt and which is truthful.

2. In the case where $|B|/d = o(n)$, the theorem allows to recover $1-o(1)$ fraction of the good nodes.

3. The proof of the theorem is algorithmic. However, the algorithm takes exponential time if we only assume that $|T| > |B|$ (or if we assume that $|T| > (1/2+\mu)n$ for a very small fixed $\mu = \mu(\lambda,d) > 0$).

The fact that the detection algorithm is not efficient even in cases when T is larger than B by a (not too large) linear number of nodes is not a coincidence. Indeed, the algorithm described in the proof of the theorem, presented in the next sections, provides a set T of more than $n/2$ truthful nodes, which is consistent with the reports obtained, when such a set exists. We show that the problem of producing such a set when it exists is NP-hard, even when restricted to bounded-degree expanders (and even if we ensure that there is such a set of size at least $n/2+\eta n$).

Theorem 1.5. *There exist $c > 0, \eta > 0$ such that the following promise problem is NP-hard. The input is a graph $G=(V,E)$ with $|V| = n$, which is an $(n,d,c\sqrt{d})$-graph along with a report ρ. The promise is that either*

- *ρ is compatible with τ satisfying $|T| = |\tau^{-1}(\mathsf{t})| \geq n/2 + \eta n$, or*
- *any τ which is compatible with ρ satisfies $|T| = \tau^{-1}(\mathsf{t}) \leq n/2 - \eta n$.*

The objective is to distinguish between the two options above.

The proof is presented in Section 2.2.

We also establish in Section 2.3 the following simple statement, which shows that at least some (weak) form of expansion is needed for solving the corruption detection problem.

Proposition 1.6. *Let $G = (V,E)$ be a graph on n nodes so that it is possible to remove at most εn nodes of G and get a graph in which each connected component is of size at most εn. Then given a report ρ compatible with an assignment τ, with $T = \tau^{-1}(\mathsf{t})$ and $B = \tau^{-1}(\mathsf{c})$, and $|T| \geq (1-2\varepsilon)n$, it is impossible to identify even a single member $t \in T$ from the reports of all nodes. In particular, this is the case for planar graphs or graphs with a fixed excluded minor even if $\varepsilon = \Theta(n^{-1/3})$.*

Note that there is still a significant gap between the expansion properties that suffice for solving the detection problem, described in Theorem 1.4, and the conditions in the last proposition that are necessary for such a solution. It will be interesting to obtain tighter relations between expansion and corruption detection. This is further discussed in Section 4.

1.2 Results for directed graphs

In this subsection we consider directed graphs (digraphs). This is motivated by the fact that in various auditing situations it is unnatural to allow u to inspect v whenever v inspects u. In fact, it may even be desirable not to allow any short cycles in the directed inspection graph.

For a set of nodes A, in a graph G, define:

$$N(A) := \{v : \exists u \in A, \{u,v\} \in E\}. \tag{1.2}$$

For a digraph $G = (V,E)$, let

$$N^+(U) = \{(v : \exists u \in U, (u,v) \in E\}, \quad N^-(U) = \{(v : \exists u \in U, (v,u) \in E\}.$$

Note that $N(U), N^+(U)$ and $N^-(U)$ may contain elements from U. We first present an explicit construction of regular directed expanders which expand at three different scales as in the undirected case.

Proposition 1.7. *There exist constants $c_1, c_2 > 0$ and infinitely many values of d for which there are infinitely many values of n and an explicit construction of $3(d-6)$-regular oriented graphs $G = ([n], E)$ which satisfy the following properties.*

- *The girth of G as an undirected graph is at least $2\log_d(n)/9$.*
- *If $A \subset [n]$ is of size at most $n/(9d)$ then $|N^+(A)| > c_1 d|A|$ and $|N^-(A)| > c_1 d|A|$.*
- *For any set $A \subset [n]$ of size at most $n/4$, it holds that $|N^+(A)| > 2|A|$ and $|N^-(A)| > 2|A|$.*

- *If $A, B \subset [n]$ with $|A| \geq c_2 n/d$, and $|B| \geq n/4$, then there is a directed edge from A to B and a directed edge from B to A.*

The proof, presented in Section 3, is based on "packing" three different group-based expanders to obtain the expansion in the different scales. Given the directed expanders constructed in Proposition 1.7, we prove the following directed analogue of Theorem 1.4.

Theorem 1.8. *Consider the oriented graphs constructed in Proposition 1.7 . There exist constants $c_3(c_1,c_2)$ and $c_4(c_1,c_2)$, for which the following holds.*

For every report ρ, there exist sets T' and B' of nodes such that for every valid truthfulness assignment τ compatible with ρ we have

$$T' \subset T, \quad B' \subset B, \quad |T \setminus T'| \leq \frac{c_3}{d}|B|, \quad |B \setminus B'| \leq \frac{c_3}{d}|B|. \tag{1.3}$$

where $T = \tau^{-1}(\mathsf{t})$ and $B = \tau^{-1}(\mathsf{c})$.

Moreover, there is a linear-time algorithm that finds subsets T' and B' satisfying (1.3) for every truthfulness assignment τ compatible with ρ which satisfies $|T| > (1 + c_4/d)(n/2)$.

1.3 Relation to distributed computing

The model considered in the paper requires a central agency that obtains the report ρ from all nodes. This is in contrast to models in distributed computing and byzantine agreement where all nodes only communicate with neighboring nodes. We note, however, that if the only objective is to ensure that most truthful nodes will know the correct types of large subsets of the set of truthful and the set of corrupt nodes, this can be achieved in the common distributed model, even if the number of truthful nodes is much less than the number of corrupt ones. This is stated in the next proposition. Its proof, which is much simpler than that of Theorem 1.4, is presented at the end of subsection 2.1.

Proposition 1.9. *Let $G = (V,E)$ be an (n,d,λ)-graph and suppose that each node of G has a unique name known to all its neighbors. Let τ be a truthful assignment with $T = \tau^{-1}(\mathsf{t})$ and $B = \tau^{-1}(\mathsf{c})$, put $|T| = tn$, and assume that $t \geq \frac{3\lambda}{d}$. Then there is an efficient distributed algorithm on the graph G in which all nodes in a subset T' of the set of truthful nodes T output subsets $T' \subset T$ and $B' \subset B$, classifying correctly the type of all nodes in these subsets, so that the following holds.*

$$\text{If } t > 1/2 \text{ then} \quad |T \setminus T'| \leq \frac{8\lambda^2}{d^2}|B|, \quad |B \setminus B'| \leq \frac{8\lambda^2}{d^2}|B|, \tag{1.4}$$

$$\text{If } \frac{3\lambda}{d} \leq t \leq 1/2 \text{ then} \quad |V \setminus (T' \cup B')| \leq \frac{3\lambda^2}{td^2}n \; (\; \leq \frac{\lambda}{d}n). \tag{1.5}$$

"Efficient" here means that every node sends and receives $O(n)$ messages.

1.4 Related work

The vast literature on corruption detection in computer science, and in particular on the diagnosable system problem and the PMC model introduced in [23], deals either with the problem of identifying all corrupt nodes, or with that of identifying a single corrupt node. As observed in [23], a necessary condition for the identification of *all* corrupt nodes in a network with t corrupt nodes is that the *minimal in-degree* in the network is at least t. Therefore, if the number of corrupt nodes is linear in the total number of nodes, all in-degrees have to be linear, and the total number of edges has to be quadratic.

The main contribution of the present paper is a proof that the number of required edges may be much smaller when relaxing the requirement of identifying *all* corrupt nodes and replacing it by the requirement of the identification of the types of most of the nodes. By relaxing the requirement as above we are able to study bounded-degree graphs, in particular d-regular graphs. Our main new result is that a linear number of edges ensures the recovery of the types of a large fraction of the nodes, provided the graph is a sufficiently strong expander. It was shown already in [23] that a linear number of edges suffices to ensure the detection of *a single* corrupt node. We show that such a small number of edges suffices to determine the types of a $1-O(1/d)$ fraction of the nodes, even when the number of truthful nodes exceeds that of corrupt ones by only 1.

In the context of Byzantine agreement it was discovered already in the 80s that expanders allow "almost everywhere agreement." This was first established by Dwork et al. in [9] and further developed in subsequent work, see, in particular [11] and its references. In [9] it is shown that on expanders one can achieve a relaxed notion of Byzantine agreement, where a large fraction of non-faulty nodes agree on a value. Like in our results, in the bounded-degree case, this fraction is bounded away from 1 unless the number of corrupt nodes is sublinear. Moreover, the results of [9] require the number of faulty nodes to be a small fraction (much smaller than $1/2$) of the total number of nodes.

It is therefore not very surprising that graph expansion is relevant to corruption detection. It is, however, interesting to see that in sufficiently strong expanders it is possible to classify most of the truthful and most of the corrupt nodes even if the number of truthful nodes exceeds the number of corrupt ones by only 1.

1.5 Techniques

The proof of Theorem 1.4 rely crucially on the fact that $(n,d,O(\sqrt{d}))$-graphs expand well at different scales.

- Every set A of size $\Omega(n/d)$ and every set B of size $n/4$ have at least one edge between them.
- Every set A of size $O(n/d)$ has $|N(A)\setminus A| \geq |A|$.
- Every set such that $|A\cup N(A)| \leq 3n/4$ satisfies $|N(A)| \geq \Omega(d)|A|$.

For the directed case, we prove the existence of directed expanders of high girth with analogous properties in Proposition 1.7.

We observe that a certain weak expansion property, namely, the nonexistence of small separators, is necessary for corruption detection. Combining this observation with the planar separator theorem of

Lipton and Tarjan and its extensions we conclude that planar graphs and graphs with a fixed excluded minor are not good for corruption detection.

Finally we discuss the algorithmic aspects of our problem using results about hardness of approximation.

2 Proofs

In this section we present the proofs of all results besides the ones for directed graphs. These appear in the next section.

2.1 Undirected graphs

We first state the following generalization of Theorem 1.4. Theorem 1.4 follows when considering graphs with $\lambda \leq 2\sqrt{d-1}$ $(< 2\sqrt{d})$.

Theorem 2.1. *Let $G = (V,E)$ be a (n,d,λ) graph, where $d^2 \geq 24\lambda^2$. Then, for every report ρ, there exist sets T' and B' of nodes such that for every valid truthfulness assignment τ compatible with ρ we have*

$$|T \setminus T'| \leq \frac{8\lambda^2}{d^2}|B|, \quad |B \setminus B'| \leq \frac{8\lambda^2}{d^2}|B|. \tag{2.1}$$

where $T = \tau^{-1}(\mathsf{t})$ and $B = \tau^{-1}(\mathsf{c})$. Moreover, there is a linear-time algorithm that finds subsets T' and B' satisfying (2.1) for every truthfulness assignment τ compatible with ρ which satisfies $|T| > (1/2 + 3\lambda^2/d^2)n$.

For a positive $\delta < 1/8$ call a graph $G = (V,E)$ on a set of n nodes a δ-*good* expander if any set U of at most $2\delta n$ nodes has more than $|U|$ neighbors outside U, and there is an edge between any pair of sets of nodes provided one of them is of size at least δn and the other is of size at least $n/4$. We recall how standard results about expanders imply that $(n,d,O(\sqrt{d}))$-graphs are δ good for $\delta = \Theta(1/d)$.

Lemma 2.2 ([2], Corollary 1)**.** *Let $G = (V,E)$ be an (n,d,λ)-graph and let $B \subset V$ be of size bn. Let tn denote the size of the set $V \setminus N(B)$, of nodes that have no neighbors in B. Then:*

$$t \leq \frac{\lambda^2}{d^2}\frac{1-b}{b}.$$

Lemma 2.3. *Any (n,d,λ)-graph in which $3\lambda^2/d^2 \leq \delta \leq 1/8$, is a δ-good expander.*

Proof. Let U be of size at least δn and suppose that that there are no edges between U and a set B of size $n/4$. Then if $u = |U|/n$, by Lemma 2.2 it follows that

$$u \leq \frac{3\lambda^2}{d^2},$$

which is a contradiction.

Similarly, let U be a set of size $un \leq 2\delta n$ and suppose that $|N(U) \setminus U| < un$. Then the set $B = V \setminus (U \cup N(U))$ is of size at least bn where $b > 1 - 2u \geq 1/2$. This implies by Lemma 2.2 that

$$u \leq \frac{\lambda^2}{d^2}\frac{1-b}{b} \leq \frac{\delta}{3}4u \leq u/2,$$

which is a contradiction. □

The other property of expanders we will need is that small sets expand by a factor of $\Omega(d)$. In particular:

Lemma 2.4. *Let $G = (V,E)$ be an (n,d,λ)-graph. Let $A \subset V$ be such that $|A \cup N(A)| \leq 3n/4$. Then*

$$|A \cup N(A)| \geq \frac{d^2}{4\lambda^2}|A|. \tag{2.2}$$

Proof. The proof is similar. Let $|A| = an$. Let $B = V \setminus (A \cup N(A))$ and put $|B| = (1-ca)n \geq n/4$. Since there are no edges between A and B, it follows that

$$a \leq \frac{\lambda^2}{d^2}\frac{ca}{1-ca} \leq 4ca\frac{\lambda^2}{d^2},$$

and so

$$c \geq \frac{d^2}{4\lambda^2},$$

which is the same as (2.2). □

Proof of Theorem 2.1. Let G be an (n,d,λ)-graph, where $d^2 \geq 24\lambda^2$.

Let τ be a truthfulness assignment consistent with ρ and let $T = \tau^{-1}(\mathsf{t})$ and $B = \tau^{-1}(\mathsf{c})$. Note that by Lemma 2.3, G is δ good, where $\delta = 3\lambda^2/d^2$.

Let H be the spanning subgraph of G in which $(u,v) \in E$ iff $\rho(u,v) = \rho(v,u) = \mathsf{t}$. Let $V_1, V_2, \ldots, V_s$ be the sets of nodes of the connected components of H.

Claim 2.5. *All the nodes of each V_i are of the same type, that is, for each $1 \leq i \leq s$, either $V_i \subset T$ or $V_i \subset B$.*

Proof. Suppose u and v are neighbors in H. If $u \in T$ then $v \in T$ (as u reports so). If $u \in B$, then $v \in B$ (as v reports that $u \in T$). □

Call a component of H truthful if it is a subset of T, otherwise it is a subset of B and we call it corrupt.

Let H' be the induced subgraph of G on the set T of all truthful nodes.

Claim 2.6. *Any connected component of H' is also a connected component of H.*

Proof. If $u, v \in T$ are adjacent in G (and hence in H'), they are adjacent in H as well, by definition and by the fact that each of them reports honestly about its neighbors. Thus each component C' of H' is contained in a component C of H. However, no $v \in T$ is adjacent in H to a node $w \in B$, implying that in fact $C' = C$ and establishing the assertion of the claim. □

Claim 2.7. *The graph H' contains a connected component of size at least $|T| - \delta n > (1/2 - \delta)n$.*

Proof. Assume this is false and the largest connected component of H' is on a set of nodes U_1 of size smaller than $|T| - \delta n$. Since the total number of nodes of H' is $|T| > n/2$, it is easy to check that one can split the connected components of H' into two disjoint sets, each of total size at least δn. However, the bigger among the two is of size bigger than $n/4$, and hence, since G is a δ-good expander, there is an edge of G between the two groups. This is impossible, as it means that there is an edge of G between two distinct connected components of H'. □

The analysis so far allow us to prove the easy part of the theorem.

Claim 2.8. *If $|T| > (1/2 + \delta)n$ then there exists a linear-time algorithm which finds $T' \subset T$ and $B' \subset B$ such that*

$$|T \setminus T'| \leq 8\frac{\lambda^2}{d^2}|B|, \quad |B \setminus B'| \leq 8\frac{\lambda^2}{d^2}|B|.$$

Proof. Note that if $|T| > (1/2 + \delta)n$ then Claim 2.7 implies that H must contain a connected component of size bigger than $n/2$, which must be truthful. Denote the nodes of this component by T' and $B' = N(T')$ and $R = V \setminus (B' \cup T')$. Clearly $B' \subset B$ and by Lemma 2.4 it follows that

$$|R \cup B'| \geq \frac{d^2}{4\lambda^2}|R|$$

which by the assumption that $d^2 \geq 24\lambda^2$ that gives $|R| + |B'| \geq 6|R|$ and hence $|R \cup B'| \leq \frac{6}{5}|B'|$ implies

$$|T \setminus T'| \leq |R| < 5\frac{\lambda^2}{d^2}|B|, \quad |B \setminus B'| \leq |R| < 5\frac{\lambda^2}{d^2}|B|,$$

establishing the required inequality (with room to spare).

Thus, if this is the case, the set T' is found by the simple, linear-time algorithm that computes the connected components of H. Furthermore, the set B' can also be found by computing the node boundary of H which is easily computed in linear time as well. □

It remains to show that even if we only assume that $|T| > n/2$ then we can still correctly classify most of the truthful and most of the corrupt nodes. We proceed with the proof of this stronger statement.

By Claims 2.6 and 2.7 it follows that H contains at least one connected component of size at least $(1/2 - \delta)n \geq 3/8n$. If H contains only one such component, then this component must consist of truthful nodes, and we can correctly classify all of them. Having these truthful nodes, we also know the types of all their neighbors. By the assumption on G this gives the types of all nodes but less than δn, thus establishing (2.1).

Otherwise, there is another connected component of size at least $(1/2 - \delta)n$, and as there is no room for more than two such components, there are exactly two of them, say V_1 and V_2. Note that by the expansion properties of G there are edges of G between V_1 and V_2, and hence it is impossible that both of them are truthful components. As one of them must be truthful, it follows that exactly one of V_1 and V_2 is a truthful component and the other corrupt. We next show that we can determine the types of both components.

Construct an auxiliary weighted graph S on the set of nodes $1,2,\ldots,s$ representing the connected components $V_1,V_2\ldots,V_s$ as follows. The weight w_i of i is defined by $w_i = |V_i|/|V|$. Two nodes i and j are connected iff there is at least one edge of G that connects a node in V_i with one in V_j. Call an independent set in the graph S *large* if its total weight is bigger than $1/2$. Note that by the discussion above T must be a union of the form $T = \bigcup_{i\in I} V_i$, where I is a large independent set in the graph S. In order to complete the argument we prove the following.

Claim 2.9. *Either there is no large independent set in S containing 1, or there is no large independent set in S containing 2.*

Proof. Assume this is false, and let I_1 be a large independent set in S containing 1, and I_2 a large independent set in S containing 2. To get a contradiction we show that for $w(I_1) = \sum_{i\in I_1} w_i$ and $w(I_2) = \sum_{i\in I_2} w_i$ we have $w(I_1)+w(I_2) \leq 1$ (and hence it is impossible that each of them has total weight bigger than a half).

To prove the above note, first, that the two nodes 1 and 2 of S are connected (as each corresponds to a set of more than $(1/2-\delta)n$ nodes of G, hence there are edges of G connecting V_1 and V_2). Therefore I_1 must contain 1 but not 2, and I_2 contains 2 but not 1.

If there are any nodes i of S connected in S both to 1 and to 2, then these nodes belong to neither I_1 nor I_2, as these are independent sets. Similarly, if a node i is connected to 1 but not to 2, then it can belong to I_2 but not to I_1, and the symmetric statement holds for nodes connected to 2 but not to 1. So far we have discussed only nodes that can belong to at most one of the two independent sets I_1 and I_2. If this is the case for all the nodes of S, then each of them contributes its weight only to one of the two sets and their total weight would thus be at most 1, implying that it cannot be that the weight of each of them is bigger than $1/2$, and completing the proof of the claim. It thus remains to deal with the nodes of S that belong to both I_1 and I_2. Let $J \subset \{3,4,\ldots,s\}$ be the set of all these nodes. Note, first, that the total weight of the nodes in J is at most 2δ, as the total weight of 1 and 2 is at least $2(1/2-\delta) = 1-2\delta$. Note also that by the discussion above each $j \in J$ is not a neighbor of 1 or of 2. By the assumption about the expander G the total weight of the nodes that are neighbors of nodes in J and do not belong to J is bigger than the total weight of the nodes in J. Indeed, this is the case as the number of neighbors in G of the set $\bigcup_{j\in J} V_j$ that do not lie in this set is bigger than the size of the set. We thus conclude that if $J' = N_S(J) - J$ denotes the set of neighbors of J that do not belong to J, then the total weight of the nodes in J' exceeds the total weight of the nodes in J, and the nodes in J' belong to neither I_1 nor I_2. We have thus proved that the sum of weights of the two independent sets I_1 and I_2 satisfies

$$w(I_1)+w(I_2) \leq 2w(J)+(1-w(J)-w(J')) \leq w(J)+w(J')+(1-w(J)-w(J')) = 1$$

contradicting the fact that both I_1 and I_2 are large. This completes the proof of the claim. □

By Claim 2.9 we conclude that one can determine the types of the components V_1 and V_2. This means that we can classify at least $(1/2-\delta)n$ truthful nodes with no error. Recall that this is the case also when H has only one connected component of size at least $(1/2-\delta)n$. Having these truthful nodes, we also know the types of all their neighbors. By the assumption on G this gives the types of all nodes but less than δn, completing the proof of (2.1).

As noted above, the algorithm described in the proof above is a linear-time algorithm provided $|T| > (1/2+\delta)n$. However, if we only assume that $|T| > |B|$ the proof provides only a non-efficient algorithm for deciding the types of the components V_1 and V_2. Indeed, we have to compute the maximum weight of an independent set containing 1 in the weighted graph S, and the maximum weight of an independent set containing 2. By the proof above, exactly one of this maxima is larger than $1/2$, providing the required types. □

The proof of Proposition 1.9 follows from the fact that it deals with the case of a large component of truthful nodes that can communicate with each other.

Proof of Proposition 1.9. For simplicity, consider a synchronized protocol. In the protocol, in each round, each node transmits to all of its neighbors the types of all nodes it recognizes as truthful and as corrupt. A truthful node adds to its record the types of all nodes it receives from truthful neighbors. Let T' denote the set of nodes of the largest connected component of the induced subgraph of G on the set T of truthful nodes and put $B' = N(T') - T'$. It is clear that all nodes in T' following the protocol classify all elements of T' as truthful and all elements of B' as corrupt. If $t > 1/2$, that is, $|T| > |B|$, then by the proof of Theorem 2.1, (1.4) holds. If $\frac{3\lambda}{d} \le t \le 1/2$ then by Lemma 2.2 T' is of size at least, say, $\frac{tn}{3}$ (since there is at least one edge between any two sets, each of size at least $\frac{tn}{3} \ge \frac{\lambda}{d}n$.) Applying Lemma 2.2 again it follows that $|V - N(T')| \le \frac{\lambda^2}{d^2}\frac{3}{t}n \le \frac{\lambda}{d}n$ and (1.5) follows as all nodes in T' classify correctly the types of all nodes in $N(T')$. □

2.2 Hardness

In this subsection we prove Theorem 1.5 which explains the non-efficiency of the algorithm in the proof of Theorem 1.4.

Proof of Theorem 1.5. The proof is based on the following result [7]: there exist constants $b < a < 1/2$ such that deciding if a graph H on m nodes, all of whose degrees are bounded by 4, has a maximum independent set of size at least $(a+b)m$ or at most $(a-b)m$ is NP-hard. It is easy to see that in fact one can assume that H is 4-regular.

Let G' be an $(n, d-4, c_1\sqrt{d})$-graph with node-set V, where $c\sqrt{d} \ge 8$. Split the nodes into 3 disjoint sets, V_1, V_2, V_3, each of size $\Omega(n)$, where V_3 is an independent set in G' of size m, all its neighbors are in V_2, $|V_1| = n/2 - am$ and $|V_2| = n/2 - m + am$, where $m = \Omega(n)$. Write η for the constant satisfying $bm = \eta n$. Add on V_3 a bounded-degree graph H as above, in which it is hard to decide if the maximum independent set is of size at least $(a+b)m$ or at most $(a-b)m$. That is, identify the set of nodes of H with V_3 and add edges between the nodes of V_3 as in H. Add a 4-regular graph on V_1 and another one on V_2. Call the resulting graph G and note that it is an $(n, d, c_1\sqrt{d}+4) = (n, d, c\sqrt{d})$ graph.

The reports of the nodes are as follows. Each node in V_1 reports true on each neighbor it has in V_1, and corrupt on any other neighbor. Similarly, each node of V_2 reports true on any neighbor it has in V_2 and corrupt on any other neighbor, and each node in V_3 reports corrupt on all its neighbors. Note that with these reports the connected components of the graph H in the proof of Theorem 1.4 are V_1, V_2 and every singleton in V_3.

It is easy to check that here if H has an independent set I of size at least $(a+b)m$, then G has a set T of truthful nodes of size at least $n/2+bm$, namely, the set $I \cup V_1$, which is consistent with all reports. If H has no independent set of size bigger than $(a-b)m$, then G does not admit any set T of truthful nodes of size bigger than $n/2-bm$ consistent with all reports. This completes the proof. □

2.3 Graphs with small separators

In this subsection we describe the simple proof of Proposition 1.6.

Proof of Proposition 1.6. Let B' be a set of at most εn nodes of G whose removal splits G into connected components with node-classes $V_1, V_2, \ldots, V_s$, each of size at most εn. Consider the following s possible scenarios R_i, for $1 \leq i \leq s$.

R_i: the set of corrupt nodes is $B = B' \cup V_i$, all the others are truthful nodes. The nodes in B' report that all their neighbors are corrupt. The nodes in V_i report that their neighbors in V_i are in T, and that all their other neighbors are in B. (The truthful nodes, of course, report truthfully about all their neighbors).

It is not difficult to check that in all these s scenarios, all nodes make exactly the same reports. On the other hand, there is no node of G that is truthful in all these scenarios, hence it is impossible to identify a truthful node with no error. Since the number of corrupt nodes in all scenarios is at most $2\varepsilon n$, the first assertion of the theorem follows. The claim regarding planar graphs and graphs with excluded minors follows from the results in [18], [5]. □

3 Directed graphs

3.1 Construction of directed expanders

Here we provide the proofs for the case of directed graphs. We start with the proof of existence of explicit directed expanders that expand in three scales. We will use the following lemma for undirected graphs from [4].

Lemma 3.1 (Lemma 3.6 [4])**.** *Let $G = (V,E)$ be an (n,d,λ) graph. Let $6/d < \gamma < 1$. Let $A, X \subset V$ be sets of nodes such that*

- $|A| \leq \frac{\gamma}{2(d+1)} n$,
- *for all $v \in A$, it holds that $|\{w \in X : (w,v) \in E\}| \geq \gamma(d+1)$.*

Then:

$$|X| \geq \frac{\gamma^2 d^2}{9\lambda^2}|A|.$$

We will also use the following variant of Lemma 2.3 whose proof is similar.

Lemma 3.2. *Any (n,d,λ)-graph in which $16\lambda^2/d^2 \leq \delta$ satisfies that any set of size at least $\delta n/2$ and any set of size $n/8$ have at least one edge between them.*

Proof. Let U be of size at least $\delta n/2$ and suppose that that there are no edges between U and a set B of size $n/8$. Then if $u = |U|/n$, by Corollary 2.2 it follows that

$$u \leq 7\frac{\lambda^2}{d^2},$$

which is a contradiction.

□

Proof of Proposition 1.7. Let $G' = ([n], E')$ be a d-regular undirected non-bipartite Ramanujan Cayley graph as constructed in [19] or [20]. This is a Cayley graph of a finite group Γ of order <– n, with respect to a set S' of d generators, $S' = \{a_1, a_1^{-1}, a_2, a_2^{-1}, \ldots, a_{d/2}, a_{d/2}^{-1}\}$. This graph is known to have girth at least $2\log_d(n)/3$, which is the same as saying there is no nontrivial word of the generators of length less than $2\log_d(n)/3$ that is the identity.

Let

$$T = \left\{a_4, a_4^{-1}, \ldots, a_{d/2}, a_{d/2}^{-1}\right\}.$$

Note that the Cayley graph of Γ with respect to T is a $(d-6)$-regular graph whose second eigenvalue is at most $2\sqrt{d}+6$. Thus if $d \geq 36$ this is an $(n, d-6, 3\sqrt{d})$-graph.

Now for $1 \leq i \leq 3$, let $T_i = a_i^{-1} T a_i$. Then clearly, G_i, the Cayley graph of Γ with respect to T_i, is also an $(n, d-6, 3\sqrt{d})$-graph. Moreover, if $S = T_1 \cup T_2 \cup T_3$, then the Cayley graph H of Γ with respect to S has girth at least

$$\frac{2}{9}\frac{\log n}{\log d},$$

since a nontrivial word in S of length k corresponds to a non-trivial word of length at most $3k$ in S'. Note also that G' is $3(d-6)$-regular.

The desired graph, $G = ([n], E)$, will be obtained by assigning directions to the edges $\bar{E}$ of H as follows:

- If $\{a, b\}$ is an edge of G_1 then orient it from a to b if $a > b$ and from b to a if $b > a$.
- If $\{a, b\}$ is an edge of G_2 then orient it from b to a if $a > b$ and from a to b if $b > a$.
- In the graph G_3 all the degrees are even. Pick an orientation of the edges of G_3 by picking a directed Eulerian cycle and orienting the edges according to it. In particular, note that every in-degree and out-degree is exactly $d/2 - 3$ in G_3.

We now verify the expansion at the three different scales. First, we apply Lemma 3.1 to the graph G_3 and sets A of size $|A| \leq n/(9d)$, and $\gamma = 1/3$ and obtain that

$$|N^+(A)| \geq \frac{(d-6)^2}{81(3\sqrt{d})^2}|A| \geq \frac{d}{400}|A|,$$

for $d \geq 36$. The proof for $N^-(A)$ is identical.

Next, let A and B be two sets of nodes with $n/16 \geq |A| \geq 200n/d$ and $|B| \geq n/4$. Let $m(A)$ and $m(B)$ denote the medians of the sets of numbers given by A and B. Then either $m(A) \geq m(B)$ or $m(B) \geq m(A)$.

If $m(A) \geq m(B)$, let $A' = \{v \in A : v \geq m(A)\}$ and $B' = \{v \in B : v \leq m(B)\}$. Then $|A'| \geq 100n/d$ and $|B'| \geq n/8$ and all elements of A' are bigger than all elements of B'. Thus by Lemma 3.2 applied to G_1, with $\delta = 200/d$, $\lambda = 3\sqrt{d}$ and $d - 6 \geq 30$, there is at least one edge in G_1 connecting A' and B'. This is a directed edge from A to B. A symmetric argument applies if $m(A) \leq m(B)$.

Note that the last statement implies that if $n/4 \geq |A| \geq 200n/d$ then the set of non-neighbors of A is of size at most $n/4$ and therefore $|N(A)| \geq 3n/4 > 2|A|$.

For sets A with $n/(10d) \leq |A| \leq 200n/d$, we have

$$|N(A)| \geq \frac{d}{400}\frac{n}{10d} = \frac{n}{4000},$$

which is greater than $400n/d$ for d sufficiently large. □

We next present the proof of Theorem 1.8, which resembles that of Theorem 1.4 but requires several additional ideas.

Proof of Theorem 1.8. Let c_1, c_2 be the constants from Proposition 1.7 and put

$$\delta = c_2/d.$$

Let τ be a truthfulness assignment consistent with ρ and let $T = \tau^{-1}(\mathsf{t})$ and $B = \tau^{-1}(\mathsf{c})$. Let H be the spanning subgraph of G in which an edge (u,v) of G is an edge of H iff $\rho(u,v) = \mathsf{t}$. Let $V_1, V_2, \ldots, V_s$ be the sets of nodes of the strongly connected components (SCCs, for short) of H.

Claim 3.3. *All the nodes of each V_i are of the same type, that is, for each $1 \leq i \leq s$, either $V_i \subset T$ or $V_i \subset B$.*

Proof. If $u \in T$ and v is an out-neighbor of u in H, then $v \in T$ (as u reports so). If $v \in B$, and u is an in-neighbor of v in H, then $u \in B$ (as u reports that $u \in T$). □

Call an SCC of H truthful if it is a subset of T, else it is a subset of B and we call it corrupt.

Let H' be the induced subgraph of G on the set T of all truthful nodes.

Claim 3.4. *Any SCC of H' is also an SCC of H.*

Proof. If $u, v \in T$ and (u,v) is an edge of G, then it is an edge of H too. Thus each SCC C' of H' is contained in an SCC C of H. This SCC is truthful, by Claim 3.3, and cannot contain any additional truthful nodes as otherwise these belong to C' as well. □

Claim 3.5. *The graph H' contains an SCC of size at least $|T| - 2\delta n > (1/2 - 2\delta)n$.*

Proof. Consider the component graph of H': this is the directed graph F whose nodes are all the SCCs of H', where there is a directed edge from C to C' iff there is some edge of H' from some node of C to some node of C'. It is easy and well known that this graph is a directed acyclic graph, and hence there is a topological order of it, that is, a numbering $C_1, C_2, \ldots, C_r$ of the components so that all edges between different components are of the form (C_i, C_j) with $i < j$. Order the nodes of H' in a linear order according to this topological order, where the nodes of C_1 come first (in an arbitrary order), those of C_2 afterwards,

etc. Let u_i be the node in place i according to this order ($1 \le i \le |T|$). If the nodes $u_{\delta n}$ and $u_{|T|-\delta n+1}$ belong to the same SCC, then this component is of size at least $|T|-2\delta n$ and we are done. Otherwise, the SCC containing $u_{|T|/2}$ differs from either that containing $u_{\delta n}$ or from that containing $u_{|T|-\delta n+1}$. In the first case, the set A of all SCCs up to that containing $u_{\delta n}$ is of size at least δn, and the set B of all SCCs starting from that containing $u_{|T|/2}$ is of size at least $|T|/2 \ge n/4$. In addition there is no edge directed from B to A, contradicting the property of G. The second case leads to a symmetric contradiction, establishing the claim. □

Note that the above shows that if $|T| > (1/2+2\delta)n$ then H' and hence also H must contain a SCC C of size bigger than $n/2$, which must be truthful. Let T' denote the set of nodes reachable from C in H and note that $T' \subset T$ and moreover $B' := N^+(T') \subset B$. Let $R = V \setminus (B' \cup T')$.

By the expansion properties of G it follows that $|R| \le c_2 n/d$ and and that $|R| \le c_3|B'|/d \le c_3|B|/d$ for a constant $c_3 = c_3(c_1,c_2)$.

We next show that even if we only assume that $|T| > n/2$ we can still classify correctly most of the truthful nodes.

By the last two claims it follows that H contains at least one SCC of size at least $(1/2-2\delta)n \ge 3/8n$. If H contains only one such component, then this component must consist of truthful nodes, and we can identify all of them (and hence also the types of all their out-neighbors). Otherwise, there is another SCC of size at least $(1/2-\delta)n$, and as there is no room for more than two such components, there are exactly two of them, say V_1 and V_2. Note that by the properties of G there are edges of G from V_1 to V_2 and from V_2 to V_1, and hence it is impossible that both of them are truthful components. As one of them must be truthful, it follows that exactly one of them is truthful and one is corrupt. We next show that we can determine the types of both components.

Recall that we have the SCCs of H, and the set T of all truthful nodes must be a union of a subset of these SCCs. In addition, this set must be of size bigger than $n/2$ and must be consistent with all reports of the nodes along every edge (in the sense that for any edge (u,v) with $u \in T$, the report of u on v should be consistent with the actual type of v).

Claim 3.6. *Given the strongly connected components $V_1, V_2, \ldots, V_s$ of H and the reports along each edge, either there is no union I_1 of SCCs including V_1 whose size exceeds $n/2$ so that $T = I_1, B = V - I_1$ is consistent with all reports along the edges, or there is no union I_2 of SCCs including V_2 whose size exceeds $n/2$ so that $T = I_2, B = V - I_2$ is consistent with all reports along the edges.*

Proof. Assume this is false, and let I_1, I_2 be as above. By the above discussion we know that I_1 contains V_1 but not V_2 and I_2 contains V_2 but not V_1. Note that if some SCC V_i is contained both in I_1 and in I_2 and there is any directed edge (u,v) from V_i to some other SCC V_j, then if the report along this edge is that v is truthful, then V_j must be truthful component in both I_1 and in I_2. Similarly, if the report along this edge is $v \in B$, then V_j must be outside I_1 and outside I_2. In particular, there are no edges at all from V_i to V_1 or V_2 (as each of them lies in exactly one of the two unions I_1, I_2). Let J be the set of all SCCs that are contained in both I_1, I_2. By the remark above, for every edge (u,v) from a node of J to a node outside J, the report along the edge must be $v \in B$ (since otherwise v would also be in an SCC which is truthful both in I_1 and in I_2 and hence would be in J). Thus all edges (u,v) as above report $v \in B$, implying that all components outside J to which there are directed edges from nodes in J belong to neither I_1 nor I_2.

By the properties of our graph the total size of these components exceeds that of J, (as $|J| \leq 4\delta n$ and all out-neighbors of J are outside V_1, V_2), and this shows that the sum of the sizes of I_1 and I_2 is at most

$$2|J| + (|V| - |J| - |N^+(J) - J|) \leq |V|.$$

Therefore it cannot be that both I_1 and I_2 are of size bigger than $|V|/2 = n/2$, proving the claim. □

By the last claim it follows that one can determine the types of the SCCs V_1 and V_2. This means that we can classify at least $(1/2 - 2\delta)n$ truthful nodes with no error. Recall that this is the case also when H has only one SCC of size at least $(1/2 - \delta)n$. Having these truthful nodes, we also know the types of all their out-neighbors. By the assumption on G this gives the types of all nodes but at most $O(|B|/d)$, completing the proof of the main part of the theorem.

The comment about the linear-time algorithm provided $|T| > (1/2 + 2\delta)n$ is clear. If we only assume that $|T| > |B|$ the proof provides only a non-efficient algorithm for deciding the types of the SCCs V_1 and V_2. Indeed, we have to check all 2^s possibilities of the types of each of the SCCs and see which ones are consistent with all reports and are of total size bigger than $n/2$. By the proof above, only one of the two SCCs V_1, V_2 will appear among the truthful SCCs of such a possibility. □

4 Discussion and open problems

The usefulness of expanders for corruption detection raises the natural question about the existence of good explicit spectral expanders with any desired number of nodes. After our work has been posted, explicit construction for such undirected expander graphs with any (large) number of nodes were obtained in [1].

It is interesting to study in more detail the relation between expansion and corruption detection.

Question 4.1. Provide sharp criteria in terms of expansion and the fractional size of the set T for enabling corruption detection.

For a weak result in the desired direction, consider the following argument. We say that an undirected graph G is *δ-connected* if for every two disjoint sets A_1, A_2 with $|A_1| \geq \delta n, |A_2| \geq (1 - 3\delta)n$ there is at least one edge between A_1 and A_2. Note that the notion of δ-connectedness is much weaker than expansion. In particular a graph G can be δ-connected, yet at the same time have $\delta n/2$ isolated nodes, while any δ-good expander must be connected.

Proposition 4.2. *Suppose that $|T| = (1 - \varepsilon)n$ and the graph G is ε-connected. Then, given the reports of all nodes, it is possible to find in linear time a subset $T' \subset T$ of size at least $(1 - 2\varepsilon)n$.*

Proof. Let $E' \subset E$ be the set of edges both of whose end-points declare each other truthful. Recall that each connected component of $G' = (V, E')$ is either truthful or corrupt.

Let $T_1, T_2, \ldots$ denote all the components of size at least εn in G'. Then we claim that for $T' = \bigcup T_i$, $|T \setminus T'| < \varepsilon n$. Assume otherwise. Since all the connected components of $T \setminus T'$ are of size at most εn, there exists $T'' \subset T \setminus T'$ of size in $[\varepsilon n, 2\varepsilon n]$ with no edges to $T \setminus T''$ whose size is in $[(1 - 3\varepsilon)n, (1 - 2\varepsilon)n]$. This is a contradiction to ε-connectedness, proving the claim. By ε-connectedness we cannot have more

than one component T_i of truthful nodes (as there are edges between any such T_i and the union of the other large components of truthful nodes), hence T' must be connected. Since it is clearly easy to find T' in linear time given the reports of all nodes, the desired result follows. □

To see that the conditions of Claim 4.2 are tight up to constant factors, consider the star graph with m leaves. Assume that $|T| \leq m-1$. Then it is easy to see that one cannot find even one member of T if all nodes declare all their neighbors corrupt. On the other hand, this example is (vacuously) $1/(4m)$ connected. To get a non-trivial example, one can replace each node with a complete graph K_k and each edge with a complete bipartite graph $K_{k,k}$ for an arbitrary k.

In follow-up work [14], the connection between expansion and corruption detection is formalized using the conjectured hardness of Small Set Expansion [24]. Assuming the hardness of Small Set Expansion, it is shown that it is computationally hard to approximate the minimal number of nodes whose corruption makes identifying even one truthful node impossible.

We conclude with a short discussion of a variant of the model. From the modeling perspective, it is interesting to consider probabilistic variants of the corruption detection problem.

Question 4.3. What is the effect of relaxing the assumption that truthful nodes always report the status correctly? Suppose for example that each truthful node reports the status of each of its neighbors independently accurately with probability $1-\varepsilon$. Note that in this case it is impossible to detect the status of an individual node with probability one. However it is still desirable to find sets T' and B' such that the symmetric difference $T \triangle T'$ and $B \triangle B'$ are small with high probability. Under what conditions can this be achieved? What are good algorithms for finding T' and B'?

Remark 4.4. Alweiss [6] addressed this problem after our work has been posted.

Acknowledgments. We thank Gireeja Ranade for suggesting to consider problems of corruption in networks, Peter Winkler for helpful comments regarding the Byzantine Agreement problem, and two anonymous referees for providing relevant references. We are also grateful to Laci Babai for numerous suggestions to improve the presentation as well as some of the results.

References

[1] NOGA ALON: Explicit expanders of every degree and size, 2020. [arXiv:2003.11673] 17

[2] NOGA ALON, JEHOSHUA BRUCK, JOSEPH SEFFI NAOR, MONI NAOR, AND RON M. ROTH: Construction of asymptotically good, low-rate error-correcting codes through pseudo-random graphs. *IEEE Trans. Inform. Theory*, 38(2):509–516, 1992. [doi:10.1109/18.119713] 8

[3] NOGA ALON AND FAN R. K. CHUNG: Explicit construction of linear sized tolerant networks. *Discrete Math.*, 72(1-3):15–19, 1988. Preliminary version in Proc. First Japan Conf. on Graph Theory and Appl., 1986. [doi:10.1016/0012-365X(88)90189-6] 3

[4] NOGA ALON, GIL KALAI, MOTY RICKLIN, AND LARRY STOCKMEYER: Lower bounds on the competitive ratio for mobile user tracking and distributed job scheduling. *Theoret. Comput. Sci.*, 130(1):175–201, 1994. Preliminary version in FOCS'92. [doi:10.1016/0304-3975(94)90158-9] 13

[5] NOGA ALON, PAUL SEYMOUR, AND ROBIN THOMAS: A separator theorem for nonplanar graphs. *J. Amer. Math. Soc.*, 3(4):801–808, 1990. [doi:10.2307/1990903] 13

[6] RYAN ALWEISS: Noisy corruption detection. *Inform. Process. Lett.*, 155(105897), 2020. [doi:10.1016/j.ipl.2019.105897, arXiv:1908.07493] 18

[7] PIOTR BERMAN AND MAREK KARPINSKI: On some tighter inapproximability results (extended abstract). In *Proc. 26th Internat. Colloq. on Automata, Languages and Programming (ICALP'99)*, volume 1644, pp. 200–209. Springer, 1999. [doi:10.1007/3-540-48523-6_17] 12

[8] THOMAS H. CORMEN, CHARLES E. LEISERSON, RONALD L. RIVEST, AND CLIFFORD STEIN: *Introduction to Algorithms*. MIT Press, Cambridge, MA, third edition, 2009. 2

[9] CYNTHIA DWORK, DAVID PELEG, NICHOLAS PIPPENGER, AND ELI UPFAL: Fault tolerance in networks of bounded degree. *SIAM J. Comput.*, 17(5):975–988, 1988. Preliminary version in STOC'86. [doi:10.1137/0217061] 7

[10] ODD-HELGE FJELDSTAD: Fighting fiscal corruption: lessons from the Tanzania Revenue Authority. *Public Administration and Development*, 23(2):165–175, 2003. [doi:10.1002/pad.278] 2

[11] JUAN A. GARAY AND RAFAIL OSTROVSKY: Almost-everywhere secure computation. In *Advances in Cryptology—EUROCRYPT'08*, volume 4965, pp. 307–323. Springer, 2008. [doi:10.1007/978-3-540-78967-3_18] 7

[12] S. LOUIS HAKIMI AND ASHOK T. AMIN: Characterization of connection assignment of diagnosable systems. *IEEE Trans. Computers*, C-23(1):86–88, 1974. [doi:10.1109/t-c.1974.223782] 2

[13] SHLOMO HOORY, NATHAN LINIAL, AND AVI WIGDERSON: Expander graphs and their applications. *Bulletin of the American Mathematical Society*, 43(4):439–561, 2006. [doi:10.1090/S0273-0979-06-01126-8] 3

[14] YAN JIN, ELCHANAN MOSSEL, AND GOVIND RAMNARAYAN: Being corrupt requires being clever, but detecting corruption doesn't. In *10th Innovations in Theoretical Comp. Sci. Conf. (ITCS'19*, volume 124, pp. 45:1–45:14. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2019. [doi:10.4230/LIPIcs.ITCS.2019.45, arXiv:1809.10325] 2, 18

[15] TIKO KAMEDA, SHUNICHI TOIDA, AND F. J. ALLAN: A diagnosing algorithm for networks. *Information and Control*, 29(2):141–148, 1975. [doi:10.1016/S0019-9958(75)90508-2] 2

[16] JON G. KUHL AND SUDHAKAR M. REDDY: Distributed fault-tolerance for large multiprocessor systems. In *Proc. 7th Ann. Symp. on Computer Architecture (ISCA'80)*, pp. 23–30. ACM Press, 1980. [doi:10.1145/800053.801905] 2

[17] LESLIE LAMPORT, ROBERT SHOSTAK, AND MARSHALL PEASE: The Byzantine generals problem. *ACM Trans. Program. Lang. Syst.*, 4(3):382–401, 1982. [doi:10.1145/357172.357176] 2

[18] RICHARD J. LIPTON AND ROBERT ENDRE TARJAN: A separator theorem for planar graphs. *SIAM J. Appl. Math.*, 36(2):177–189, 1979. [doi:10.1137/0136016] 13

[19] ALEX LUBOTZKY, RALPH PHILLIPS, AND PETER SARNAK: Ramanujan graphs. *Combinatorica*, 8(3):261–277, 1988. [doi:10.1007/BF02126799] 3, 14

[20] GRIGORY A. MARGULIS: Explicit group-theoretic constructions of combinatorial schemes and their applications in the construction of expanders and concentrators (Russian). *Problemy Peredachi Informatsii*, 24(1):51–60, 1988. Link at Mathnet.ru. English translation: Problems Inform. Transmission 24 (1988/1) 39–46. 3, 14

[21] BISWANATH MUKHERJEE, L. TODD HEBERLEIN, AND KARL N. LEVITT: Network intrusion detection. *IEEE Network*, 8(3):26–41, 1994. [doi:10.1109/65.283931] 2

[22] RICHARD P. NIELSEN: Corruption networks and implications for ethical corruption reform. *Journal of Business Ethics*, 42(2):125–149, 2003. [doi:10.1023/A:1021969204875] 2

[23] FRANCO P. PREPARATA, GERNOT METZE, AND ROBERT T. CHIEN: On the connection assignment problem of diagnosable systems. *IEEE Trans. Electronic Computers*, EC-16(6):848–854, 1967. [doi:10.1109/PGEC.1967.264748] 2, 7

[24] PRASAD RAGHAVENDRA AND DAVID STEURER: Graph expansion and the unique games conjecture. In *Proc. 42nd STOC*, pp. 755–764. ACM Press, 2010. [doi:10.1145/1806689.1806792] 18

[25] MICHAEL T. ROCK AND HEIDI BONNETT: The comparative politics of corruption: accounting for the East Asian paradox in empirical studies of corruption, growth and investment. *World Development*, 32(6):999–1017, 2004. [doi:10.1016/j.worlddev.2003.12.002] 2

[26] MAŁGORZATA STEINDER AND ADARSHPAL S. SETHI: A survey of fault localization techniques in computer networks. *Sci. Comput. Programming*, 53(2):165–194, 2004. [doi:10.1016/j.scico.2004.01.010] 2

AUTHORS

Noga Alon
Professor
Department of Mathematics
Princeton University
and
Schools of Mathematics and Computer Science
Tel Aviv University
nogaa@tau.ac.il

Elchanan Mossel
Professor
Department of Mathematics and IDSS, 2-434
Massachusetts Institute for Technology
Cambridge, MA 02139, USA
elmos@mit.edu

Robin Pemantle
Professor
Department of Mathematics
University of Pennsylvania
209 South 33rd Street
Philadelphia, PA 19104, USA
pemantle@math.upenn.edu

ABOUT THE AUTHORS

NOGA ALON is a Professor of Mathematics at Princeton University and a Professor Emeritus of Mathematics and Computer Science at Tel Aviv University, Israel. He received his Ph. D. in Mathematics at the Hebrew University of Jerusalem in 1983 and held visiting positions in various research institutions, including MIT, the Institute for Advanced Study in Princeton and Microsoft Research Redmond and Herzliya.

He is an ACM Fellow and an AMS Fellow, a member of the Israel Academy of Sciences and Humanities, of the Academia Europaea and of the Hungarian Academy of Sciences, and received the Erdős Prize, the Feher Prize, the Pólya Prize, the Bruno Memorial Award, the Landau Prize, the Gödel Prize, the Israel Prize, the EMET Prize, the Dijkstra Prize and the Nerode Prize.

His research interests are mainly in Combinatorics, Graph Theory and their applications in Theoretical Computer Science. His main contributions include the study of expander graphs and their applications, the investigation of derandomization techniques, the foundation of streaming algorithms, the development and applications of algebraic and probabilistic methods in Discrete Mathematics and the study of problems in Information Theory, Combinatorial Geometry and Combinatorial Number Theory.

ELCHANAN MOSSEL is a professor of mathematics at the Massachusetts Institute of Technology. His primary research fields are probability theory, combinatorics, and statistical inference. He received his Ph. D. in Mathematics at the Hebrew University of Jerusalem in 2000. He held a postdoctoral position at Microsoft Research, Redmond and was a Miller Research Fellow at UC Berkeley before becoming a Professor at UC Berkeley, the Weizmann Institute, the University of Pennsylvania and finally MIT.

Mossel's research spans a number of topics across mathematics, statistics, economics, and computer science, including combinatorial statistics, discrete function inequalities, isoperimetry, game theory, social choice, computational complexity, and computational evolutionary biology.

Mossel held a Sloan Fellowship. He is a fellow of the American Mathematical Society and a Simons Fellow.

ROBIN PEMANTLE grew up in Berkeley in the 1960s and 70s. He became interested in probability theory when Persi Diaconis, a mathematician and former magician, visited M.I.T. He was supposed to be working on a dissertation in another subject, but found himself working mostly on problems he heard from Persi. During a year off traveling in the South Pacific, he ended up working on probability theory in various youth hostels and on boats. Pemantle got his Ph. D. in 1988, spent three years on post-doctoral fellowships at Berkeley, Cornell, and Oregon State, most of the 1990s at UW-Madison, three years at Ohio State, and is now at Penn.

Pemantle's research focuses on two areas. Within Probability Theory, the research concerns discrete probability models, including random graph theory, processes with reinforcement, statistical models and random walks. The other research area, analytic combinatorics, is the subject of a textbook with Mark C. Wilson (2013).

Pemantle has also been interested in Mathematics Education from an early age, growing up with an insider's view of an alternative school, and teaching mathematics to grades 5-8 during his college years and before.

Pemantle has held a Sloan Fellowship, a Presidential Faculty Fellowship, the Rollo Davidson Prize, and a Lilly Teaching Fellowship. He was a top five finisher in the Putnam Competition and is a Simons Fellow, a fellow of the American Mathematical Society and a fellow of the Institute for Mathematical Statistics.