

Robustly Self-Ordered Graphs: Constructions and Applications to Property Testing

Oded Goldreich Avi Wigderson

May 20, 2022

Abstract

A graph G is called *self-ordered* (a.k.a asymmetric) if the identity permutation is its only automorphism. Equivalently, there is a unique isomorphism from G to any graph that is isomorphic to G . We say that $G = (V, E)$ is *robustly self-ordered* if the size of the symmetric difference between E and the edge-set of the graph obtained by permuting V using any permutation $\pi : V \rightarrow V$ is proportional to the number of non-fixed-points of π . In this work, we initiate the study of the structure, construction and utility of robustly self-ordered graphs.

We show that robustly self-ordered bounded-degree graphs exist (in abundance), and that they can be constructed efficiently, in a strong sense. Specifically, given the index of a vertex in such a graph, it is possible to find all its neighbors in polynomial-time (i.e., in time that is poly-logarithmic in the size of the graph).

We provide two very different constructions, in tools and structure. The first, a direct construction, is based on proving a sufficient condition for robust self-ordering, which requires that an auxiliary graph is expanding. The second construction is iterative, boosting the property of robust self-ordering from smaller to larger graphs. Structurally, the first construction always yields expanding graphs, while the second construction may produce graphs that have many tiny (sub-logarithmic) connected components.

We also consider graphs of unbounded degree, seeking correspondingly unbounded robustness parameters. We again demonstrate that such graphs (of linear degree) exist (in abundance), and that they can be constructed efficiently, in a strong sense. This turns out to require very different tools. Specifically, we show that the construction of such graphs reduces to the construction of non-malleable two-source extractors (with very weak parameters but with some additional natural features).

We demonstrate that robustly self-ordered bounded-degree graphs are useful towards obtaining lower bounds on the query complexity of testing graph properties both in the bounded-degree and the dense graph models. Indeed, their robustness offers efficient, local and distance preserving reductions from testing problems on ordered structures (like sequences) to the unordered (effectively unlabeled) graphs. One of the results that we obtain, via such a reduction, is a subexponential separation between the query complexities of testing and tolerant testing of graph properties in the bounded-degree graph model.

⁰The authors' affiliation and grant acknowledgements appear in the Acknowledgements section.

Contents

1	Introduction	1
1.1	Robustly self-ordered bounded-degree graphs	2
1.1.1	Our main results and motivation	2
1.1.2	Techniques	3
1.2	Robustly self-ordered dense graphs	6
1.2.1	Our main results	6
1.2.2	Techniques	7
1.3	Perspective	8
1.4	Roadmaps	8
 Part I The Case of Bounded-Degree Graphs		 9
2	The Edge-Colored Variant	10
2.1	Transformation to standard (uncolored) version	11
2.2	Application: Making the graph regular and expanding	15
2.3	Local computability of the transformations	16
3	The Direct Construction	18
3.1	A sufficient condition for robust self-ordering of directed colored graphs	19
3.2	From the directed variant to the undirected one	22
4	The Three-Step Construction	24
4.1	Existence	24
4.2	Constructions	27
4.3	Strong (i.e., local) constructions	34
4.4	Local self-ordering	35
5	Application to Testing Bounded-Degree Graph Properties	38
5.1	Background	39
5.2	Our first result and the general methodology	40
5.3	Separating tolerant testing from testing	45
6	Random Regular Graphs are Robustly Self-Ordered	46
 Part II The Case of Dense Graphs		 51
7	Existence and Transformation to Bounded-Degree Graphs	51
8	Relation to Non-Malleable Two-Source Extractors	55
8.1	The first construction	56
8.2	The second construction	62
8.3	Obtaining efficient self-ordering	68
9	Application to Testing Dense Graph Properties	70
10	The Case of Intermediate Degree Bounds	73
Acknowledgements		78
References		78

1 Introduction

For a (labeled) graph $G = (V, E)$, and a bijection $\phi : V \rightarrow V'$, we denote by $\phi(G)$ the graph $G' = (V', E')$ such that $E' = \{\{\phi(u), \phi(v)\} : \{u, v\} \in E\}$, and say that G' is isomorphic to G . The set of automorphisms of the graph $G = (V, E)$, denoted $\mathbf{aut}(G)$, is the set of permutations that preserve the graph G ; that is, $\pi \in \mathbf{aut}(G)$ if and only if $\pi(G) = G$. We say that a graph is asymmetric (equiv., self-ordered) if its set of automorphisms is a singleton, which consists of the trivial automorphism (i.e., the identity permutation). We actually prefer the term *self-ordered*, because we take the perspective that is offered by the following equivalent definition.

Definition 1.1 (self-ordered (a.k.a asymmetric) graphs): *The graph $G = ([n], E)$ is self-ordered if for every graph $G' = (V', E')$ that is isomorphic to G there exists a unique bijection $\phi : V' \rightarrow [n]$ such that $\phi(G') = G$.*

In other words, given an isomorphic copy $G' = (V', E')$ of a fixed graph $G = ([n], E)$, there is a unique bijection $\phi : V' \rightarrow [n]$ that orders the vertices of G' such that the resulting graph (i.e., $\phi(G')$) is identical to G . Indeed, if $G' = G$, then this unique bijection is the identity permutation.¹

In this work, we consider a feature, which we call *robust self-ordering*, that is a quantitative version self-ordering. Loosely speaking, a graph $G = ([n], E)$ is robustly self-ordered if, for every permutation $\pi : [n] \rightarrow [n]$, the size of the symmetric difference between G and $\pi(G)$ is proportional to the number of non-fixed-points under π ; that is, $|E \Delta \{\{\pi(u), \pi(v)\} : \{u, v\} \in E\}|$ is proportional to $|\{i \in [n] : \pi(i) \neq i\}|$. (In contrast, self-ordering only means that the size of the symmetric difference is positive if the number of non-fixed-points is positive.)

Definition 1.2 (robustly self-ordered graphs): *A graph $G = (V, E)$ is said to be γ -robustly self-ordered if for every permutation $\pi : V \rightarrow V$ it holds that*

$$|E \Delta \{\{\pi(u), \pi(v)\} : \{u, v\} \in E\}| \geq \gamma \cdot |\{i \in [n] : \pi(i) \neq i\}|, \quad (1)$$

where Δ denotes the symmetric difference operation. An infinite family of graphs $\{G_n = ([n], E_n)\}_{n \in \mathbb{N}}$ (such that each G_n has maximum degree d) is called **robustly self-ordered** if there exists a constant $\gamma > 0$, called the **robustness parameter**, such that for every n the graph G_n is γ -robustly self-ordered.

Note that $|E_n \Delta \{\{\pi(u), \pi(v)\} : \{u, v\} \in E_n\}| \leq 2d \cdot |\{i \in [n] : \pi(i) \neq i\}|$ always holds (for families of maximum degree d). The term “robust” is inspired by the property testing literature (cf. [32]), where it indicates that some “parametrized violation” is reflected proportionally in some “detection parameter”.

The second part of Definition 1.2 is tailored for bounded-degree graphs, which will be our focus in Section 2–6. Nevertheless, in Sections 7–10 we consider graphs of unbounded degree and unbounded robustness parameters. In this case, for a function $\rho : \mathbb{N} \rightarrow \mathbb{R}$, we say that an infinite family of graphs $\{G_n = ([n], E_n)\}_{n \in \mathbb{N}}$ is ρ -robustly self-ordered if for every n the graph G_n is $\rho(n)$ -robustly self-ordered. Naturally, in this case, the graphs must have $\Omega(\rho(n) \cdot n)$ edges.² In Sections 7–9 we consider the case of $\rho(n) = \Omega(n)$.

¹Naturally, we are interested in efficient algorithms that find this unique ordering, whenever it exists; such algorithms are known when the degree of the graph is bounded [30].

²Actually, all but at most one vertex must have degree at least $\rho(n)/2$.

1.1 Robustly self-ordered bounded-degree graphs

The first part of this paper (i.e., Section 2–6) focuses on the study of robustly self-ordered bounded-degree graphs.

1.1.1 Our main results and motivation

We show that robustly self-ordered (n -vertex) graphs of bounded-degree not only exist (for all $n \in \mathbb{N}$), but can be efficiently constructed in a strong (or local) sense. Specifically, we prove the following result.

Theorem 1.3 (constructing robustly self-ordered bounded-degree graphs): *For all sufficiently large $d \in \mathbb{N}$, there exist an infinite family of d -regular robustly self-ordered graphs $\{G_n\}_{n \in \mathbb{N}}$ and a polynomial-time algorithm that, given $n \in \mathbb{N}$ and a vertex $v \in [n]$ in the n -vertex graph G_n , finds all neighbors of v (in G_n).*

We stress that the algorithm runs in time that is polynomial in the description of the vertex; that is, the algorithm runs in time that is polylogarithmic in the size of the graph. Theorem 1.3 holds both for graphs that consists of connected components of logarithmic size and for “strongly connected” graphs (i.e., expanders).

Recall that given an isomorphic copy G' of such a graph G_n , the original graph G_n (i.e., along with its unique ordering) can be found in polynomial-time [30]. Furthermore, we show that the pre-image of each vertex of G' in the graph G_n (i.e., its index in the aforementioned ordering) can be found in time that is polylogarithmic in the size of the graph (see discussion in Section 4.4, culminating in Theorem 4.7).³

We present two proofs of Theorem 1.3. Loosely speaking, the first proof reduces to proving that a $2d$ -regular n -vertex graph representing the action of d permutations on $[n]$ is robustly self-ordered if the $2 \cdot \binom{n}{2}$ -vertex graph representing the action of these permutations on (ordered) vertex-pairs is an expander.⁴ The graphs constructed in this proof are expanders, whereas the graphs constructed via by the second proof can be either expanders or consist of connected components of logarithmic size. More importantly, the graphs constructed in the second proof are coupled with *local self-ordering and local reversed self-ordering algorithms* (see Section 4.4). The second proof proceeds in three steps, starting from the mere existence of robustly self-ordered bounded-degree ℓ -vertex graphs, which yields a construction that runs in $\text{poly}(\ell^\ell)$ -time. Next, a $\text{poly}(n)$ -time construction of n -vertex graphs is obtained by using the former graphs as small subgraphs (of $o(\log n)$ -size). Lastly, strong (a.k.a local) constructability is obtained in an analogous manner. For more details, see Section 1.1.2.

We demonstrate that robustly self-ordered bounded-degree graphs are useful towards obtaining lower bounds on the query complexity of testing graph properties in the bounded-degree graph model. Specifically, we use these graphs as a key ingredient in a general methodology of transporting lower bounds regarding testing binary strings to lower bounds regarding testing graph properties

³The algorithm asserted above is said to perform *local self-ordering* of G' according to G_n . For $\phi(G') = G_n$, given a vertex v in G' , this algorithm returns $\phi(v)$ in $\text{poly}(\log n)$ -time. In contrast, a *local reversed self-ordering* algorithm is given a vertex $i \in [n]$ of G_n and returns $\phi^{-1}(i)$. The second algorithm is also presented in Section 4.4 (see Theorem 4.9).

⁴Here and throughout this paper, by *expander* we mean families of bounded-degree graphs that have constant expansion (cf., e.g., [26]).

in the bounded-degree graph model. In particular, using the methodology, we prove the following two results.

1. A *subexponential separation between the query complexities of testing and tolerant testing of graph properties in the bounded-degree graph model*; that is, for some constant $c > 0$, the query complexity of tolerant testing is at least $\exp(q^c)$, where q is the query complexity of standard testing.

This result, which appears as Theorem 5.5, is obtained by transporting an analogous result that was known for testing binary strings [15].

2. A linear query complexity lower bound for testing an *efficiently recognizable* graph property in the bounded-degree graph model, where the lower bound holds even if the tested graph is restricted to consist of connected components of logarithmic size (see Theorem 5.2).

As discussed in Section 5, an analogous result was known in the general case (i.e., without the restriction on the size of the connected components), and we consider it interesting that the result holds also in the special case of graphs with small connected components.

To get a feeling of why robustly self-ordered graphs are relevant to such transportation, recall that strings are ordered objects, whereas graphs properties are effectively sets of unlabeled graphs, which are unordered objects. Hence, we need to make the graphs (in the property) ordered, and furthermore make this ordering robust in the very sense that is reflected in Definition 1.2. We comment that the theme of reducing ordered structures to unordered structures occurs often in the theory of computation and in logic, and is often coupled with analogues of query complexity.

Lastly, in Section 6, we prove that random $2d$ -regular graphs are robustly self-ordered; see Theorem 6.1. This extends work in probabilistic graph theory, which proves a similar result for the weaker notion of self-ordering [4, 5].

1.1.2 Techniques

As stated above, we present two different constructions that establish Theorem 1.3: A direct construction and a three-step construction. Both constructions utilize a variant of the notion of robust self-ordering that refers to edge-colored graphs, which we review first.

The edge-coloring methodology. At several different points, we found it useful to start by demonstrating the robust self-ordering feature in a relaxed model in which edges are assigned a constant number of colors, and the symmetric difference between graphs accounts also for edges that have different colors in the two graphs (see Definition 2.1). This allows us to analyze different sets of edges separately.

For example, we actually analyze the direct construction in the edge-colored model, while associating each of the underlying d permutations with a different color. This association allows for analyzing the effect of each permutation separately (see below). Another example, which arises in the three-step construction, occurs when we super-impose a robustly self-ordered graph with an expander graph in order to make the robustly self-ordered graph expanding (as needed for the second and third step of the aforementioned three-step construction). In this case, assigning the edges of each of the two graphs a different color, allows for easily retaining the robust self-ordering feature (of the first graph).

We obtain robustly self-ordered graphs (in the original sense) by replacing all edges that are assigned a specific color with copies of a constant-sized (asymmetric) gadget, where different (and in fact non-isomorphic) gadgets are used for different edge colors. The soundness of this transformation is proved in Theorem 2.4.

The direct construction. For any d permutations, $\pi_1, \dots, \pi_d : [n] \rightarrow [n]$, we consider the *Schreier graph* (see [26, Sec. 11.1.2]) defined by the action of these permutation on $[n]$; that is, the edge-set of this graph is $\{\{v, \pi_i(v)\} : v \in [n] \& i \in [d]\}$. Loosely speaking, we prove that *this $2d$ -regular n -vertex graph is robustly self-ordered if another Schreier graph is an expander*. The second Schreier graph represents the action of the same permutations on *pairs* of vertices (in $[n]$); that is, this graph consisting of the vertex-set $\{(u, v) : u, v \in [n]\}$ and the edge-set $\{\{(u, v), (\pi_i(u), \pi_i(v))\} : u, v \in [n] \& i \in [d]\}$.⁵

The argument is actually made with respect to edge-colored directed graphs (i.e., the edge-set of the first graph is $\{(v, \pi_i(v)) : v \in [n] \& i \in [d]\}$ and the directed edge $(v, \pi_i(v))$ is assigned the color i). Hence, we also present a transformation of robustly self-ordered edge-colored directed graphs to analogous undirected graphs. Specifically, we replace the directed edge (u, v) colored j by a 2-path with a designated auxiliary vertex $a_{u,v,j}$, while coloring the edge $\{u, a_{u,v,j}\}$ by $2j - 1$ and the edge $\{a_{u,v,j}, v\}$ by $2j$.

We comment that permutations satisfying the foregoing condition can be efficiently constructed; for example, any set of expanding generators for $\text{SL}_2(p)$ (e.g., the one used by [29]) yield such permutations on $[n] \equiv \{(1, i) : i \in \text{GF}(p)\} \cup \{(0, 1)\}$ (see Proposition 3.3).⁶

The three-step construction. Our alternative construction of robustly self-ordered (bounded-degree) n -vertex graphs proceeds in three steps.

1. First, we prove the existence of bounded-degree n -vertex graphs that are robustly self-ordered (see Theorem 4.1), while observing that this yields a $\exp(O(n \log n))$ -time algorithm for constructing them.
2. Next (see Theorem 4.2), we use the latter algorithm to construct robustly self-ordered n -vertex bounded-degree graphs that consist of 2ℓ -sized connected components, where $\ell = \frac{O(\log n)}{\log \log n}$; these connected components are far from being isomorphic to one another, and are constructed using robustly self-ordered ℓ -vertex graphs as a building block. This yields an algorithm that constructs the n -vertex graph in $\text{poly}(n)$ -time, since $\exp(O(\ell \log \ell)) = \text{poly}(n)$.
3. Lastly, we derive Theorem 1.3 (restated as Theorem 4.5) by repeating the same strategy as in Step 2, but using the construction of Theorem 4.2 for the construction of the small connected components (and setting $\ell = O(\log n)$). This yields an algorithm that finds the neighbors of a vertex in the n -vertex graph in $\text{poly}(\log n)$ -time, since $\text{poly}(\ell) = \text{poly}(\log n)$.

The foregoing description of Steps 2 and 3 yields graphs that consists of small connected components. We obtain analogous results for “strongly connected” graphs (i.e., expanders) by superimposing these graphs with expander graphs (while distinguishing the two types of edges by using

⁵Equivalently, we consider only pairs of distinct vertices; that is, the vertex-set $\{(u, v) : u, v \in [n] \& u \neq v\}$.

⁶In this case, the primary Schreier graph represents the natural action of the group on the 1-dimensional subspaces of $\text{GF}(p)^2$.

colors (see the foregoing discussion)). In fact, it is essential to perform this transformation (on the result of Step 2) before taking Step 3; the transformation itself appears in the proof of Theorem 2.6.

Using large collections of pairwise far apart permutations. One ingredient in the foregoing three-step construction is the use of a single ℓ -vertex robustly self-ordered (bounded-degree) graph towards obtaining a *large* collection of 2ℓ -vertex (bounded-degree) graphs such that every two graphs are far from being isomorphic to one another, where “large” means $\exp(\Omega(\ell \log \ell))$ in one case (i.e., in the proof of Theorem 4.2) and $\exp(\Omega(\ell))$ in another case (i.e., in the proof of Theorem 4.5). Essentially, this is done by constructing a large collection of permutations of $[\ell]$ that are pairwise far-apart, and letting the i^{th} graph consists of two copies of the ℓ -vertex graph that are matched according to the i^{th} permutation (see the aforementioned proofs). (Actually, we use two robustly self-ordered ℓ -vertex graphs that are far from being isomorphic (e.g., have different degree).)

A collection of $L = \exp(\Omega(\ell \log \ell))$ pairwise far-apart permutations over $[\ell]$ can be constructed in $\text{poly}(L)$ -time by selecting the permutations one by one, while relying on the existence of a permutation that augments the current sequence (while preserving the distance condition, see the proof of Theorem 4.2). A collection of $L = \exp(\Omega(\ell))$ pairwise far-apart permutations over $[\ell]$ can be locally constructed such that the i^{th} permutation is constructed in $\text{poly}(\ell)$ -time by using sequences of disjoint transpositions determined via a good error correcting code (see the proof of Theorem 4.5).

The foregoing discussion begs the challenge of obtaining a construction of a collection of $L = \exp(\Omega(\ell \log \ell))$ permutations over $[\ell]$ that are pairwise far-apart along with a polynomial-time algorithm that, on input $i \in [L]$, returns a description of the i^{th} permutation (i.e., the algorithm should run in $\text{poly}(\log L)$ -time). We meet this challenge in [23]. Note that such a collection constitutes a an asymptotically good code over the alphabet $[\ell]$, where the permutations are the codewords (being far-apart corresponds to constant relative distance and $\log L = \Omega(\log(\ell!))$ corresponds to constant rate).

On the failure of some natural approaches. We mention that natural candidates for robustly self-ordered bounded-degree graphs fail. In particular, there exist expander graphs that are not robustly self-ordered. In fact, any Cayley graph is symmetric (i.e., has non-trivial automorphisms).⁷

In light of the above, it is interesting that expansion *can* serve as a sufficient condition for robust self-ordering (as explained in the foregoing review of the direct construction); recall, however, that this works for Schreier graphs, and expansion needs to hold for the action on vertex-pairs.

On optimization: We made no attempt to minimize the degree bound and maximize the robustness parameter. Note that we can obtain 3-regular robustly self-ordered graphs by applying degree reduction; that is, given a d -regular graph, we replace each vertex by a d -cycle and use each of these vertices to “hook” one original edge. To facilitate the analysis, we may use one color for the edges of the d -cycles and another color for the other (i.e., original) edges.⁸ Hence, the issue at hand is actually one of maximizing the robustness parameter of the resulting 3-regular graphs.

⁷Specifically, multiplying the vertex labels (say, on the right) by any non-zero group element yields a non-trivial automorphism (assuming that edges are defined by multiplying with a generator on the left). Such automorphisms cannot be constructed in general for Schreier graphs, and some Schreier graphs have no automorphisms (e.g., the ones we construct here).

⁸Needless to say, we later replace all colored edges by copies of adequate (3-regular) constant-sized gadgets.

Caveat (tedious): Whenever we assert a d -regular n -vertex graph, we assume that the trivial conditions hold; specifically, we assume that $n > d$ and that nd is even (or, alternatively, allow for one exceptional vertex of degree $d - 1$).

1.2 Robustly self-ordered dense graphs

In the second part of this paper (i.e., Sections 7–10) we consider graphs of unbounded degree, seeking correspondingly unbounded robustness parameters. In particular, we are interested in n -vertex graphs that are $\Omega(n)$ -robustly self-ordered, which means that they must have $\Omega(n^2)$ edges.

The construction of $\Omega(n)$ -robustly self-ordered graphs offers yet another alternative approach towards the construction of bounded-degree graphs that are $\Omega(1)$ -robustly self-ordered. Specifically, we show that n -vertex graphs that are $\Omega(n)$ -robustly self-ordered can be efficiently transformed into $O(n^2)$ -vertex bounded-degree graphs that are $\Omega(1)$ -robustly self-ordered; see Proposition 7.2, which is essentially proved by the “degree reduction via expanders” technique, while using a different color for the expanders’ edges, and then using gadgets to replace colored edges (see Theorem 2.4).

1.2.1 Our main results

It is quite easy to show that random n -vertex graphs are $\Omega(n)$ -robustly self-ordered (see Proposition 7.1); in fact, the proof is easier than the proof of the analogous result for bounded-degree graphs (Theorem 6.1). Unfortunately, constructing n -vertex graphs that are $\Omega(n)$ -robustly self-ordered seems to be no easier than constructing robustly self-ordered bounded-degree graphs. In particular, it seems to require completely different techniques and tools.

Theorem 1.4 (constructing $\Omega(n)$ -robustly self-ordered graphs): *There exist an infinite family of dense $\Omega(n)$ -robustly self-ordered graphs $\{G_n\}_{n \in \mathbb{N}}$ and a polynomial-time algorithm that, given $n \in \mathbb{N}$ and a pair of vertices $u, v \in [n]$ in the n -vertex graph G_n , determines whether or not u is adjacent to v in G_n .*

Unlike in the case of bounded-degree graphs, in general, we cannot rely on an efficient isomorphism test for finding the original ordering of G_n , when given an isomorphic copy of it. However, we can obtain dense $\Omega(n)$ -robustly self-ordered graphs for which this ordering can be found efficiently (see Theorem 8.10).

Our proof of Theorem 1.4 is by a reduction to the construction of non-malleable two-source extractors, where a suitable construction of the latter was provided by Chattopadhyay, Goyal, and Li [7]. We actually present two different reductions (Theorems 8.3 and 8.7), one simpler than the other but yielding a less efficient construction when combined with the known constructions of extractors. We mention that the first reduction (Theorem 8.3) is partially reversible (see Proposition 8.5, which reverses a special case captured in Remark 8.4).

We show that $\Omega(n)$ -robustly self-ordered n -vertex graphs can be used to transport lower bounds regarding testing binary strings to lower bounds regarding testing graph properties in the dense graph model. This general methodology, presented in Section 9, is analogous to the methodology for the bounded-degree graph model, which is presented in Section 5.

We mention that in a follow-up work [24], we employed this methodology in order to resolve several open problems regarding the relation between adaptive and non-adaptive testers in the dense graph model. In particular, we proved that there exist graph properties for which any non-adaptive tester must have query complexity that is almost quadratic in the query complexity of

the best general (i.e., adaptive) tester, whereas it has been known for a couple of decades that the query complexity of non-adaptive testers is at most quadratic in the query complexity of adaptive testers.

The case of intermediate degree bounds. Lastly, in Section 10, we consider n -vertex graphs of degree bound $d(n)$, for every $d : \mathbb{N} \rightarrow \mathbb{N}$ such that $d(n) \in [\Omega(1), n]$. Indeed, the bounded-degree case (studied in Section 2–6) and the dense graph case (studied in Sections 7–9) are special cases (which correspond to $d(n) = O(1)$ and $d(n) = n$). Using results from these two special cases, we show how to construct $\Omega(d(n))$ -robustly self-ordered n -vertex graphs of maximum degree $d(n)$, for all $d : \mathbb{N} \rightarrow \mathbb{N}$.

1.2.2 Techniques

As evident from the foregoing description, we reduce the construction of $\Omega(n)$ -robustly self-ordered n -vertex graphs to the construction of non-malleable two-source extractors.

Non-malleable two-source extractors were introduced in [8], as a variant on seeded (one-source) non-malleable extractors, which were introduced in [12]. Loosely speaking, we say that $\text{nmE} : \{0, 1\}^\ell \times \{0, 1\}^\ell \rightarrow \{0, 1\}^m$ is a non-malleable two-source extractor for a class of sources \mathcal{C} if for every two independent sources in \mathcal{C} , denoted X and Y , and for every two functions $f, g : \{0, 1\}^\ell \rightarrow \{0, 1\}^\ell$ that have no fixed-point it holds that $(\text{nmE}(X, Y), \text{nmE}(f(X), g(Y)))$ is close to $(U_m, \text{nmE}(f(X), g(Y)))$, where U_m denotes the uniform distribution over $\{0, 1\}^m$. We show that a non-malleable two-source extractor for the class of ℓ -bit sources of min-entropy $\ell - O(1)$, with a single output bit (i.e., $m = 1$) and constant error, suffices for constructing $\Omega(n)$ -robustly self-ordered n -vertex graphs. Recall that constructions with much stronger parameters (e.g., min-entropy $\ell - \ell^{\Omega(1)}$, negligible error, and $m = \ell^{\Omega(1)}$) were provided by Chattopadhyay, Goyal, and Li [7, Thm. 1]. (These constructions are quite complex. Interestingly, we are not aware of a simpler way of obtaining the weaker parameters that we need.)

Actually, we show two reductions of the construction of $\Omega(n)$ -robustly self-ordered n -vertex graphs to the construction of non-malleable two-source extractors. In both cases we use extractors that operate on pairs of sources of length $\ell = \log_2 n - O(1)$ that have min-entropy $k = \ell - O(1)$, hereafter called (ℓ, k) -sources. The extractor is used to define a bipartite graph with 2^ℓ vertices on each side, and a clique is placed on the vertices of one side so that a permutation that maps vertices from one side to the other side yields a proportional symmetric difference (between the original graph and the resulting graph).

The first reduction, presented in Theorem 8.3, requires the extractor to be *quasi-orthogonal*, which means that the residual functions obtained by any two different fixings of one of the extractor’s two arguments are almost unbiased and uncorrelated. Using the fact that non-malleable two-source extractors for (ℓ, k) -sources can be made quasi-orthogonal in $\exp(\ell)$ -time, we obtain an explicit construction of $\Omega(n)$ -robustly self-ordered n -vertex graphs (i.e., the n -vertex graph is constructed in $\text{poly}(n)$ -time).

The second reduction, presented in Theorem 8.7, yields a strongly explicit construction as asserted in Theorem 1.4 (i.e., the adjacency predicate of the n -vertex graph is computable in $\text{poly}(\log n)$ -time). This reduction uses an arbitrary non-malleable two-source extractor, and shifts the quasi-orthogonality condition to two auxiliary bipartite graphs.

Both reductions are based on the observation that if the number of non-fixed-points (of the permutation) is very large, then the non-malleability condition implies a large symmetric difference

(between the original graph and the resulting graph). This holds as long as there are at least $\Omega(2^\ell)$ non-fixed-points on each of the two sides of the corresponding bipartite graph (which corresponds to the extractor). The complementary case is handled by the quasi-orthogonality condition, and this is where the two reductions differ.

The simpler case, presented in the first construction (i.e., Theorem 8.3), is that the extractor itself is quasi-orthogonal. In this case we consider the non-fixed-points on the side that has more of them. The quasi-orthogonality condition gives us a contribution of approximately $0.5 \cdot 2^\ell$ units per each non-fixed-point, whereas the upper-bound on the number of non-fixed-points on the other side implies that most of these contributions actually count in the symmetric difference (between the original graph and the resulting graph).

In the second construction (i.e., Theorem 8.7), we augment the foregoing 2^ℓ -by- 2^ℓ bipartite graph, which is now determined by any non-malleable extractor, with an additional $4 \cdot 2^\ell$ -vertex clique that is connected to the two original 2^ℓ -vertex sets by a bipartite graph that is merely quasi-orthogonal. The analysis is analogous to the one used in the proof of Theorem 8.3, but is slightly more complex because we are dealing with a slightly more complex graph.

Errata regarding the original posting. We retract the claims made in our initial posting [22] regarding the construction of non-malleable two-source extractors (which are quasi-orthogonal) as well as the claims about the construction of relocation-detecting codes (see Theorems 1.5 and 1.6 in the original version).⁹ The source of trouble is a fundamental flaw in the proof of [22, Lem. 9.7], which may as well be wrong.

1.3 Perspective

Asymmetric graphs were famously studied by Erdos and Renyi [14], who considered the (absolute) distance of asymmetric graphs from being symmetric (i.e., the number of edges that should be removed or added to a graph to make it symmetric), calling this quantity the *degree of asymmetry*. They studied the extremal question of determining the largest possible degree of asymmetry of n -vertex graphs (as a function of n). We avoided the term “robust asymmetry” because it could be confused with the degree of asymmetry, which is a very different notion. In particular, the degree of asymmetry cannot exceed twice the degree of the graph (e.g., by disconnecting two vertices), whereas our focus is on robustly self-ordered graphs of bounded-degree.

We mention that Bollobas proved that, *for every constant $d \geq 3$, almost all d -regular graphs are asymmetric* [4, 5]. This result was extended to varying $d \in [3, n - 4]$ by Kim, Sudakov, and Vu [27]. We also mention that their proof of [27, Thm. 3.1] implies that a random n -vertex Erdos–Renyi graph with edge probability p is $2p(1 - p)n$ -robustly self-ordered.

1.4 Roadmaps

This work consists of two parts. The first part (Sections 2–6) refers to bounded-degree graphs, and the second part (Sections 7–10) refers to dense graphs. These parts are practically independent of one another, except that Theorem 10.3 builds upon Section 6. Even when focusing on one of these two parts, its contents may attract attention from diverse perspectives. Each such perspective may benefit from a different roadmap.

⁹In [22] quasi-orthogonality is called niceness; we prefer the current term, which is less generic.

Efficient combinatorial constructions. As mentioned above, in the regime of bounded-degree graphs we present two different constructions that establish Theorem 1.3. Both constructions make use of the edge-colored model and the transformations presented in Section 2. The direct construction is presented in Section 3, and the three-step construction appears in Section 4. The three-step construction is augmented by local self-ordering and local reversed self-ordering algorithms (see Section 4.4).¹⁰ In the regime of dense graphs, Sections 7 and 8 refer to the constructability of a couple of combinatorial objects; see roadmap “for the dense case” below.

Potential applications to property testing. In Section 5 we demonstrate applications of Theorem 1.3 to proving lower bounds (on the query complexity) for the bounded-degree graph testing model. Specifically, we present a methodology of transporting bounds regarding testing properties of strings to bounds regarding testing properties of bounded-degree graphs. The specific applications presented in Section 5 rely on Section 4. For the first application (Theorem 5.2) the construction presented in Section 4.2 suffices; for the second application (i.e., Theorem 5.5, which establishes a separation between testing and tolerant testing in the bounded-degree graph model), the local computation tasks studied in Section 4.4 are needed. An analogous methodology for the dense graph testing model is presented in Section 9.

Properties of random graphs. As stated above, it turns out that random $O(1)$ -regular graphs are robustly self-ordered. This result is presented in Section 6, and this section can be read independently of any other section. (In addition, Section 7 presents a proof that random (dense) n -vertex graphs are $O(n)$ -robustly self-ordered.)

The dense case and non-malleable two-source extractors. The regime of dense graphs is studied in Sections 7–9, where the construction of such graphs is undertaken in Section 8. In Section 7, we show that $\Omega(n)$ -robustly self-ordered n -vertex graphs provide yet another way of obtaining $\Omega(1)$ -robustly self-ordered bounded-degree graphs. In Section 8, we reduce the construction of $O(n)$ -robustly self-ordered n -vertex graphs to the construction of non-malleable two-source extractors. As outlined in Section 1.2.2, we actually present two different reductions, where a key issue is the quasi-orthogonality condition.

Lastly, in Section 10, for every $d : \mathbb{N} \rightarrow \mathbb{N}$ such that $d(n) \in [\Omega(1), n]$, we show how to construct n -vertex graphs of maximum degree $d(n)$ that are $\Omega(d(n))$ -robustly self-ordered. Some of the results and techniques presented in this section are also relevant to the setting of bounded-degree graphs.

Part I

The Case of Bounded-Degree Graphs

As stated in Section 1.1.2, a notion of robust self-ordering of edge-colored graphs plays a pivotal role in our study of robustly self-ordered bounded-degree graphs. This notion as well as a transformation from it to the uncolored version (of Definition 1.2) is presented in Section 2.

¹⁰For a locally constructable G_n and $G' = \phi^{-1}(G_n)$, a *local self-ordering* algorithm is given a vertex v in G' , and returns $\phi(v)$. In contrast, a *local reversed self-ordering* algorithm is given a vertex $i \in [n]$ of G_n and returns $\phi^{-1}(i)$. Both algorithms run in $\text{poly}(\log n)$ -time.

In Section 3, we present a direct construction of $O(1)$ -regular robustly self-ordered edge-colored graphs; applying the foregoing transformation, this provides our first proof of Theorem 1.3. Our second proof of Theorem 1.3 is presented in Section 4, and consists of a three-step process (as outlined in Section 1.1.2). Sections 3 and 4 can be read independently of one another, but both rely on Section 2.

In Section 5 we demonstrate the applicability of robustly self-ordered bounded-degree graphs to property testing; specifically, to proving lower bounds (on the query complexity) for the bounded-degree graph testing model. For these applications, the global notion of constructability, established in Section 4.2, suffices. This construction should be preferred over the direct construction presented in Section 3, because it can also yield graphs with small connected components. More importantly, the subexponential separation between the complexities of testing and tolerant testing of graph properties (i.e., Theorem 5.5) relies on the construction of Section 4 and specifically on the local computation tasks studied in Section 4.4.

Lastly, in Section 6, we prove that random $O(1)$ -regular graphs are robustly self-ordered. This section may be read independently of any other section.

2 The Edge-Colored Variant

Many of our arguments are easier to make in a model of (bounded-degree) graphs in which edges are colored (by a bounded number of colors), and where one counts the number of mismatches between colored edges. Namely, an edge that appears in one (edge-colored) graph contributes to the count if it either does not appear in the other (edge-colored) graph or appears in it under a different color. Hence, we define a notion of robust self-ordering for edge-colored graphs. We shall then transform robustly self-ordered edge-colored graphs to robustly self-ordered ordinary (uncolored) graphs, while preserving the degree, the asymptotic number of vertices, and other features such as expansion and degree-regularity. Specifically, the transformation consists of replacing the colored edges by copies of different connected, asymmetric (constant-sized) gadgets such that different colors are reflected by different gadgets.

We start by providing the definition of the edge-colored model. Actually, for greater flexibility, we will consider multi-graphs; that is, graphs with possible parallel edges and self-loops. Hence, we shall consider multi-graphs $G = (V, E)$ coupled with an edge-coloring function $\chi: E \rightarrow \mathbb{N}$, where E is a multi-set containing both pairs of vertices and singletons (representing self-loops). Actually, it will be more convenient to represent self-loops as 2-element multi-sets containing two copies of the same vertex.

Definition 2.1 (robust self-ordering of edge-colored multi-graphs): *Let $G = (V, E)$ be a multi-graph with colored edges, where $\chi: E \rightarrow \mathbb{N}$ denotes this coloring, and let E_i denote the multi-set of edges colored i (i.e., $E_i = \{e \in E : \chi(e) = i\}$). We say that (G, χ) is γ -robustly self-ordered if for every permutation $\mu: V \rightarrow V$ it holds that*

$$\sum_{i \in \mathbb{N}} \left| E_i \Delta \{ \{ \mu(u), \mu(v) \} : \{u, v\} \in E_i \} \right| \geq \gamma \cdot |\{i \in V : \mu(i) \neq i\}|, \quad (2)$$

where $A \Delta B$ denotes the symmetric difference between the multi-sets A and B ; that is $A \Delta B$ contains t occurrences of e if the absolute difference between the number of occurrences of e in A and B equals t .

(Definition 1.2 is obtained as a special case when the multi-graph is actually a graph and all edges are assigned the same color.)

We stress that whenever we consider “edge-colored graphs” we actually refer to edge-colored multi-graphs (i.e., we explicitly allow parallel edges and self-loops).¹¹ In contrast, whenever we consider (uncolored) graph, we refer to simple graphs (with no parallel edges and no self-loops).

Our transformation of robustly self-ordered edge-colored multi-graphs to robustly self-ordered ordinary graphs depends on the number of colors used by the multi-graph. In particular, γ -robustness of edge-colored multi-graph that uses c colors gets translated to $(\gamma/f(c))$ -robustness of the resulting graph, where $f : \mathbb{N} \rightarrow \mathbb{N}$ is an unbounded function. Hence, we focus on coloring functions that use a constant number of colors, denoted c . That is, fixing a constant $c \in \mathbb{N}$, we shall consider multi-graphs $G = (V, E)$ coupled with an edge-coloring function $\chi : E \rightarrow [c]$.

2.1 Transformation to standard (uncolored) version

As a preliminary step for the transformation, we add self-loops to all vertices and make sure that parallel edges are assigned different colors. The self-loops make it easy to distinguish the original vertices from auxiliary vertices that are parts of gadgets introduced in the main transformation. Different colors assigned to parallel edges are essential to the mere asymmetry of the resulting graph, since we are going to replace edges of the same color by copies of the same gadget.

Construction 2.2 (preliminary step towards Construction 2.3): *For a fixed $d \geq 3$, given a multi-graph $G = (V, E)$ of maximum degree d and an edge-coloring function $\chi : E \rightarrow [c]$, we define a multi-graph $G = (V, E')$ and an edge-coloring function $\chi' : E' \rightarrow [d \cdot c + 1]$ as follows.*

1. *For every pair of vertices u and v that are connected by few parallel edges, denoted $e_{u,v}^{(1)}, \dots, e_{u,v}^{(d')}$, we change, for each $i \in [d']$, the color of $e_{u,v}^{(i)}$ to $\chi'(e_{u,v}^{(i)}) \leftarrow (i - 1) \cdot c + \chi(e_{u,v}^{(i)})$. This includes also the case $u = v$.*
2. *We augment the multi-graph with self-loops colored $d \cdot c + 1$; that is, E' is the multi-set $E \cup \{e_v : v \in V\}$, where e_v is a self-loop added to v , and $\chi'(e_v) = dc + 1$.*

(Other edges $e \in E$ maintain their color; that is, for them $\chi'(e) = \chi(e)$ holds).

(For simplicity, we re-color all parallel edges, save the first one, rather than re-coloring only parallel edges that have the same color.) Note that refining the coloring may only increase the robustness parameter of an edge-colored multi-graph. Clearly, G' preserves many features of G . In particular, it preserves γ -robust self-ordering, expansion, degree-regularity, and the number of vertices.

As stated above, our transformation of edge-colored multi-graphs to ordinary graphs uses gadgets, which are constant-size graphs. Specifically, when handling a multi-graph of maximum degree d with edges that are colored by c colors, we shall use c different *connected and asymmetric* graphs.

¹¹We comment that a seemingly more appealing definition can be used for edge-colored (simple) graphs. Specifically, in that case (i.e., $E \subseteq \binom{V}{2}$), we can extend $\chi : E \rightarrow \mathbb{N}$ to non-edges by defining $\chi(\{u, v\}) = 0$ if $\{u, v\} \notin E$, and say that (G, χ) is γ -robustly self-ordered if for every permutation $\mu : V \rightarrow V$ it holds that

$$\left| \left\{ \{u, v\} \in \binom{V}{2} : \chi(\{\mu(u), \mu(v)\}) \neq \chi(\{u, v\}) \right\} \right| \geq \gamma \cdot |\{i \in V : \mu(i) \neq i\}|.$$

Furthermore, in order to maintain d -regularity, we shall use d -regular graphs as gadgets; and in order to have better control on the number of vertices in the resulting graph, each of these gadgets will contain $k = k(d, c)$ vertices. The existence of such (d -regular) asymmetric (and connected) graphs is well-known, let alone that it is known that a random d -regular k -vertex graph is asymmetric (for any constant $d \geq 3$) [4, 5].

We stress that the different gadgets are each connected and asymmetric, and it follows that they are not isomorphic to one another. We designate in each gadget an edge $\{p, q\}$, called the **designated edge**, such that omitting this edge does not disconnect the gadget. The endpoints of this edge will be used to connect two vertices of the original multi-graph. Specifically, we replace each edge $\{u, v\}$ (of the original multi-graph) that is colored i by a copy of the i^{th} gadget, while omitting its designated edge $\{p, q\}$, and connecting u to p and v to q . The construction is spelled out below.

We say that a (non-simple) multi-graph $G = (V, E)$ coupled with an edge-coloring χ is eligible if each of its vertices contains a self-loop, and parallel edges are assigned different colors. Recall that eligibility comes almost for free (by applying Construction 2.2). We shall apply the following construction only to eligible edge-colored multi-graphs.

Construction 2.3 (the main transformation): *For a fixed $d \geq 3$ and c , let $k = k(d, c)$ and G_1, \dots, G_c be different asymmetric and connected d -regular graphs over the vertex-set $[k]$. Given a multi-graph $G = (V, E)$ of maximum degree d and an edge-coloring function $\chi : E \rightarrow [c]$, we construct a graph $G' = (V', E')$ as follows.*

Suppose that the multi-set E has size m . Then, for each $j \in [m]$, if the j^{th} edge of E connects vertices u and v , and is colored i , then we replace it by a copy of G_i , while omitting its designated edge and connecting one of its endpoints to u and the other endpoint to v .

Specifically, assuming that $V = [n]$ and recalling that j is the index of the edge (colored i) that connects u and v , let $G_i^{u,v}$ be an isomorphic copy of G_i that uses the vertex set $\{n + (j - 1) \cdot k + i : i \in [k]\}$. Let $\{p, q\}$ be the designated edge in $G_i^{u,v}$, and $\hat{G}_i^{u,v}$ be the graph that results from $G_i^{u,v}$ by omitting $\{p, q\}$. Then, we replace the edge $\{u, v\}$ by $\hat{G}_i^{u,v}$, and add the edges $\{u, p\}$ and $\{v, q\}$.

Hence, $V' = [n + m \cdot k]$ and E' consists of the edges of all $\hat{G}_i^{u,v}$'s as well as the edges connecting the endpoint of the corresponding designated edges to the corresponding vertices u and v .

We stress that, although G may have parallel edges and self-loops, the graph G' has neither parallel edges nor self-loops. Also note that G' preserve various properties of G such as degree-regularity, number of connected components, and expansion (up to a constant factor).

We shall show that if the edge-colored multi-graph $G = (V, E)$ is robustly self-ordered (in the edge-colored sense), then the resulting graph $G' = (V', E')$ is robustly self-ordered (in the ordinary sense). The proof of this fact relies on a correspondence between the colored edges of G and the gadgets in G' . For starters, suppose that the permutation $\mu' : V' \rightarrow V'$ maps V to V (i.e., $\mu'(V) = V$), and gadgets to the corresponding gadgets; that is, if μ' maps the vertex-pair $(u, v) \in V^2$ to $(\mu'(u), \mu'(v)) \in V^2$, then μ' maps the vertices in the possible gadget that connects u and v to the vertices in the gadget that connects $\mu'(u)$ and $\mu'(v)$. In such a case, letting μ be the restriction of μ' to V , a difference of D colored edges between G and $\mu(G)$ translates to a difference

of at least D edges between G' and $\mu'(G')$, due to the difference between the gadgets that replace the corresponding (colored) edges of G' , whereas the number of non-fixed-point vertices in μ' is k times larger than the number of non-fixed-point vertices in μ . Assuming that G is γ -robustly self-ordered, it follows that μ has at most D/γ non-fixed-points. Hence, in this case we have

$$\frac{|G' \Delta \mu'(G')|}{|\{v \in V' : \mu'(v) \neq v\}|} \geq \frac{D}{k \cdot |\{v \in V : \mu(v) \neq v\}|} \geq \frac{D}{k \cdot D/\gamma}$$

which equals γ/k . However, in general, μ' needs not satisfy the foregoing condition. Nevertheless, if μ' splits some gadget or maps some gadget in a manner that is inconsistent with the vertices of V connected by it, then this gadget contributes at least one unit to the difference between G' and $\mu'(G')$, whereas the number of non-fixed-point vertices in this gadget is at most k . Lastly, if μ' maps vertices of a gadget to other vertices in the same gadget, then we get a contribution of at least one unit due to the asymmetry of the gadget. The foregoing argument is made rigorous in the proof of the following theorem.

Theorem 2.4 (from edge-colored robustness to standard robustness): *For constant $d \geq 3$ and $c \in \mathbb{N}$, suppose that the multi-graph $G = (V, E)$ coupled with $\chi : E \rightarrow [c]$ is eligible and γ -robustly self-ordered. Then, the graph $G' = (V', E')$ resulting from Construction 2.3 is $(\gamma/3k)$ -robustly self-ordered, where $k = k(d, c) > d$ is the number of vertices in a gadget (as determined above) and $\gamma \leq 1$.*

Proof: As a warm-up, let us verify that G' is asymmetric. We first observe that the vertices of G are uniquely identified (in G'), since they are the only vertices that are incident at copies of the gadget that replaces the self-loops.¹² Hence, any automorphism of G' must map V to V . Consequently, for any i , such an automorphism μ' must map each copy of G_i to a copy of G_i , which means that when permuting V according to μ' the edges of G as well as their colors are preserved. By the ‘‘colored asymmetry’’ of G , this implies that μ' maps each $v \in V$ to itself, and consequently each copy of G_i must be mapped (by μ') to itself. Finally, using the asymmetry of the G_i ’s, it follows that each vertex of each copy of G_i is mapped to itself. Hence, μ' must be the identity permutation.

We now turn to proving that G' is actually robustly self-ordered. Considering an arbitrary permutation $\mu' : V' \rightarrow V'$, we lower-bound the distance between G' and $\mu'(G')$ as a function of the number of non-fixed-points under μ' (i.e., of $v \in V'$ such that $\mu'(v) \neq v$). We do so by considering the contribution of each non-fixed-point to the distance between G' and $\mu'(G')$. We first recall the fact that the vertices of V (resp., of gadgets) are uniquely identified in $\mu'(G')$ by virtue of the gadgets that replace self-loops (see the foregoing warm-up).

Case 1: *Vertices of some copy of G_i that are not mapped by μ' to a single copy of G_i ; that is, vertices in some $G_i^{u,v}$ that are not mapped by μ' to some $G_i^{u',v'}$.*

(This includes the case of vertices w' and w'' of some $G_i^{u,v}$ such that $\mu'(w')$ is in $G_{i'}^{u',v'}$ and $\mu'(w'')$ is in $G_{i''}^{u'',v''}$, but $(i', u', v') \neq (i'', u'', v'')$. It also includes the case of a copy of G_i that is mapped by μ' to a copy of G_j for $j \neq i$, and the case that a vertex w in some $G_i^{u,v}$ that is mapped by μ' to a vertex in V .)

¹²Indeed, a vertex of G' is in V if and only if omitting it from G' yields several connected components such that (at least) one of them is a copy of the gadget that replaces the self-loops (with the designated edge missing).

The set of vertices $S_i^{u,v}$ of each such copy (i.e., $G_i^{u,v}$) contribute at least one unit to the difference between G' and $\mu'(G')$, since $\mu'(S_i^{u,v})$ induces a copy of \hat{G}_i in $\mu(G')$ but not in G' , where here we also use the fact that the \hat{G}_i 's are connected (and not isomorphic (for the case of $i' = i'' \neq i$)). Note that the total contribution of all vertices of the current case equals at least the number of gadgets in which they reside. Hence, if the current case contains n_1 vertices, then their contribution to the distance between G' and $\mu'(G')$ is at least n_1/k .

Ditto for vertices that do not belong to a single copy of G_i and are mapped by μ' to a single copy of G_i . This also includes $v \in V$ being mapped to some copy of some G_i , but in this case we get a contribution of one unit (rather than $1/k$ amortized units) per each such vertex (i.e., $v \in V$ such that $\mu(v) \notin V$).

Case 2: Vertices of some copy of G_i that are mapped by μ' to a single copy of G_i , while not preserving their indices inside G_i .

(This refers to vertices of some $G_i^{u,v}$ that are mapped by μ' to vertices of $G_i^{u',v'}$, where (u', v') may but need not equal (u, v) , such that for some $j \in [k]$ the j^{th} vertex of $G_i^{u,v}$ is not mapped by μ' to the j^{th} vertex of $G_i^{u',v'}$.)¹³

By the fact that G_i is asymmetric, it follows that each such copy contributes at least one unit to the difference between G' and $\mu'(G')$, and so (again) the total contribution of all these vertices is proportional to their number; that is, if the number of vertices in this case is n_2 , then their contribution is at least n_2/k .

Case 3: Vertices $v \in V$ such that $\mu'(v) \in V \setminus \{v\}$.

(This is the main case, and here we use the hypothesis that the edge-colored multi-graph G is robustly self-ordered.

Intuitively, the hypothesis that the edge-colored G is robustly self-ordered implies that such vertices contribute proportionally to the difference between the colored versions of the multi-graphs G and $\mu(G)$, where μ is the restriction of μ' to V . Indeed, we first assume, for simplicity, that $\mu'(V) = V$, an assumption we shall have to dispose of later. In this case, the number of tuples $(\{u, w\}, i)$ such that $\{u, w\}$ is colored i in exactly one of these multi-graph (i.e., either in G or in $\mu(G)$ but not in both) is at least $\gamma \cdot |\{v \in V : \mu(v) \neq v\}|$. Assuming, without loss of generality that $\chi(\{u, w\}) = i$ but either $\{\mu^{-1}(u), \mu^{-1}(w)\} \notin E$ or $\chi(\{\mu^{-1}(u), \mu^{-1}(w)\}) = j \neq i$, we observe that $\mu^{-1}(u)$ and $\mu^{-1}(w)$ cannot be connected in G' via a copy of G_i . We consider two sub-cases:

1. μ' maps a copy of G_i to $G_i^{u,w}$, but either $\mu^{-1}(u)$ or $\mu^{-1}(w)$ is not connected to this copy in G' . In this sub-case we get a contribution of at least one unit, since u and w are connected to $G_i^{u,w}$ in G' .
2. μ' does not map a copy of G_i to $G_i^{u,w}$. In this sub-case, it follows that some vertices that do not belong to a copy of G_i are mapped by μ' to $G_i^{u,w}$ which means that Case 1 applies for each such a tuple.

¹³Recall that $G_i^{u,v}$ and $G_i^{u',v'}$ are both copies of the k -vertex graph G_i , which is an asymmetric graph, and so the notion of the j^{th} vertex in them is well-defined. Formally, the j^{th} vertex of $G_i^{u,v}$ is $\phi^{-1}(j)$ such that ϕ is the (unique) bijection satisfying $\phi(G_i^{u,v}) = G_i$.

Hence, if the number of vertices in the current case is n_3 , then the number of tuples (handled by the two sub-cases) is at least $\gamma \cdot n_3$, and we get a contribution of at least $\gamma \cdot n_3/k$ (since the second sub-case is handled via Case 1).

The foregoing description is based on the assumption that $\mu(V) = V$. If this does not hold, then we redefine μ such that $\mu(v) \notin V$ is modified such that $\mu(v) = r$ if $r \in V$ has no preimage under μ' . (Of course, each such r is only used once.) Indeed, the modified μ may be fictitiously charged with d edges per each modification, but each such modification arises due to $v \in V$ that contributes at least one unit via (the last part of) Case 1. Hence, the amortized over-counting of $d \cdot \gamma/k$ units is compensated by the unit contributed in Case 1.

Case 4: *Vertices of some copy of G_i that are mapped by μ' to a different copy of G_i .*

This refers to the case that μ' maps $G_i^{u,v}$ to $G_i^{u',v'}$ such that $(u', v') \neq (u, v)$, which corresponds to mapping the gadget to a gadget connecting a different pair of vertices (but by an edge of the same color).

For u, v, u', v' and i as above, if $\mu'(u) = u'$ and $\mu'(v) = v'$, then a gadget that connects u and v in G' is mapped to a gadget that does not connect them in $\mu'(G')$ (but rather connects the vertices u' and v' , whereas either $u' \neq u$ or $v' \neq v$). So, due to the gadget-edge incident at either u or v , we get a contribution of at least one unit to the difference between G' and $\mu'(G')$, whereas the number of vertices in this gadget is k . Hence, the contribution is proportional to the number of non-fixed-points of the current type. Otherwise (i.e., $(\mu'(u), \mu'(v)) \neq (u', v')$), we get a vertex as in Case 3, and get a proportional contribution again.

Hence, the contribution of each of these cases to the difference between G' and $\mu'(G')$ is proportional to the number of vertices involved. Specifically, if there are n_i vertices in Case i , then we get a contribution-count of at least $\gamma \cdot \sum_{i \in [4]} n_i/k$, where some of these contributions were possibly counted thrice. The claim follows. \blacksquare

Remark 2.5 (fitting any desired number of vertices): *Assuming that the hypothesis of Theorem 2.4 can be met for any sufficiently large $n \in S \subseteq \mathbb{N}$, Construction 2.3 yields robustly self-ordered n' -vertex graphs for any $n' \in \{k \cdot n : n \in S\}$, where $k = k(d, c)$ is as in Theorem 2.4. To obtain such graphs also for n' that is not a multiple of k , we may use two gadgets with a different number of vertices for replacing at least one of the sets of colored edges.*

2.2 Application: Making the graph regular and expanding

We view the edge-colored model as an intermediate locus in a two-step methodology for constructing robustly self-ordered graphs of bounded-degree. First, one constructs edge-colored multi-graphs that are robustly self-ordered in the sense of Definition 2.1, and then converts them to ordinary robustly self-ordered graphs (in the sense of Definition 1.2), by using Construction 2.3 (while relying on Theorem 2.4).

We demonstrate the usefulness of this methodology by showing that it yields a simple way of making robustly self-ordered graphs be also expanding as well as regular, while maintaining a bounded degree. We just augment the original graph by super-imposing an expander (on the same vertex set), while using one color for the edges of the original graph and another color for the edges of the expander. Note that we do not have to worry about the possibility of creating parallel edges

(since they are assigned different colors). The same method applies in order to make the graph regular. We combine both transformations in the following result, which we shall use in the sequel.

Theorem 2.6 (making the graph regular and expanding): *For constant $d \geq 3$ and γ , there exists an efficient algorithm that given a γ -robustly self-ordered graph $G = (V, E)$ of maximum degree d , returns a $(d + O(1))$ -regular multi-graph expander coupled with a 2-coloring of its edges such that the edge-colored multi-graph is γ -robustly self-ordered (in the sense of Definition 2.1).*

The same idea can be applied to edge-colored multi-graphs; in this case, we use one color more than given. We could have avoided the creation of parallel edges with the same color by using more colors, but preferred to relegate this task to Construction 2.2, while recalling that it preserves both the expansion and the degree-regularity. Either way, applying Theorem 2.4 to the resulting edge-colored multi-graph, we obtain robustly self-ordered (uncolored) graphs.

Proof: For any $d'' \geq d + d'$, given a graph $G = (V, E)$ of maximum degree d that is γ -robustly self-ordered and a d' -regular expander graph $G' = (V, E')$, we construct the desired d'' -regular multi-graph G'' by super-imposing the two graphs on the same vertex set, while assigning the edges of each of these graphs a different color. In addition, we add edges to make the graph regular, and color them using the same color as used for the expander.¹⁴ Details follow.

- We superimpose G and G' (i.e., create a multi-graph $(V, E \cup E')$), while coloring the edges of G (resp., G') with color 1 (resp., color 2).

Note that this may create parallel edges, but with different colors.

- Let $d_v \leq d + d'$ denote the degree of vertex v in the resulting multi-graph. Then, we add edges to this multi-graph so that each vertex has degree d'' . These edges will also be colored 2.

(Here, unless we are a bit careful, we may introduce parallel edges that are assigned the same color. This can be avoided by using more colors for these added edges, but in light of Construction 2.2 (which does essentially the same) there is no reason to worry about this aspect.)

(Recall that the resulting edge-colored multi-graph is denoted G'' .)

The crucial observation is that, since the edges of G are given a distinct color in G'' , the added edges do not harm the robust self-ordering feature of G . Hence, for any permutation $\mu : V \rightarrow V$, any vertex-pair that contributes to the symmetric difference between G and $\mu(G)$, also contributes to an inequality between colored edges of G'' and $\mu(G'')$ (by virtue of the edges colored 1). ■

2.3 Local computability of the transformations

In this subsection, we merely point out that the transformation presented in Constructions 2.2 and 2.3 as well as the one underlying the proof of Theorem 2.6 preserve efficient local computability (e.g., one can determine the neighborhood of a vertex in the resulting multi-graph by making a polylogarithmic number of neighbor-queries to the original multi-graph). Actually, this holds provided that we augment the (local) representation of graphs, in a natural manner.

¹⁴We assume for simplicity that $|V'|$ is even. Alternatively, assuming that G contains no isolated vertex, we first augment it with an isolated vertex and apply the transformation on the resulting graph. Yet another alternative is to consider only even d'' .

Recall that the standard representation of bounded-degree graphs is by their incidence functions. Specifically, a graph $G = ([n], E)$ of maximum degree d is represented by the incident function $g : [n] \times [d] \rightarrow [n] \cup \{0\}$ such that $g(v, i) = u \in [n]$ if u is the i^{th} neighbor of v , and $g(v, i) = 0$ if v has less than i neighbors. This does not allow us to determine the identity of the j^{th} edge in G , nor even to determine the number of edges in G , by making a polylogarithmic number of queries to g , whereas this determination is needed for a local implementation of Construction 2.3. Nevertheless, efficient local computability is preserved if we use the following local representation (presented for edge-colored multi-graphs).

Definition 2.7 (local representation): *For $d, c \in \mathbb{N}$, a local representation of a multi-graph $G = ([n], E)$ of maximum degree d that is coupled with a coloring $\chi : E \rightarrow [c]$ is provided by the following three functions:*

1. *An incidence function $g_1 : [n] \times [d] \rightarrow \mathbb{N} \cup \{0\}$ such that $g_1(v, i) = j \in \mathbb{N}$ if j is the index of the i^{th} edge that incident at vertex v , and $g_1(v, i) = 0$ if v has less than i incident edges.*
2. *An edge enumeration function $g_2 : \mathbb{N} \rightarrow (\binom{[n]}{2} \times [c]) \cup \{0\}$ such that $g_2(j) = (\{u, v\}, \chi(e_j))$ if the j^{th} edge, denoted e_j , connects the vertices u and v , and $g_2(j) = 0$ if the multi-graph has less than j edges.*
3. *An vertex enumeration (by degree) function $g_3 : [d] \times [n] \rightarrow [n] \cup \{0\}$ such that $g_3(i, j) = v \in [n]$ if v is the j^{th} vertex of degree i in the multi-graph, and $g_3(i, j) = 0$ if the multi-graph has less than j vertices of degree i .*

The aforementioned incident function $g : [n] \times [d] \rightarrow [n] \cup \{0\}$ can be computed by composing g_1 and g_2 ; in particular, $g(v, i) = u \in [n]$ if $g_2(g_1(v, i)) = (\{u, v\}, k)$ for some $k \in [c]$. Needless to say, the function g_3 is redundant in the case that we are guaranteed that the multi-graph is regular. One may augment the foregoing representation by providing also the total number of edges, but this number can be determined by binary search.

Theorem 2.8 (the foregoing transformations preserve local computability): *The local representation of the multi-graph that result from Construction 2.2 can be computed by making a polylogarithmic number of queries to the given multi-graph. The same holds for Construction 2.3 and for the transformation underlying the proof of Theorem 2.6.*

Proof: For Construction 2.2, we mostly need to enumerate all parallel edges that connect u and v . This can be done easily by querying the incidence function on $(u, 1), \dots, (u, d)$ and querying the edge enumeration function on the non-zero answers. (In addition, when adding a self-loop on vertex $v \in [n]$, we need to determine the degree of v as well as the number of edges in the multi-graph (in order to know how to index the self-loop in the incidence and edge enumeration functions, respectively).)

For Construction 2.3, we merely need to determine the color of the j^{th} edge, its endpoint (and its index in the incidence list of each of its endpoints), in order to replace this colored edge by the relevant gadget. Recall that the relevant gadget uses the vertices $n + (j - 1) \cdot k + 1, \dots, n + j \cdot k$ and its edges are determined by the color of the edge that it replaces.

For the transformation underlying the proof of Theorem 2.6, adding edges to make the multi-graph regular requires determining the index of a vertex in the list of all vertices of the same degree

(in order to properly index the added edges). Here is where we use the vertex enumeration (by degree) function. (We also need a local procedure I for transforming a sorted n -long sequence $(d_1, \dots, d_n) \in [d'']$ into an all- d'' sequence by making pairs of increments; that is, given $j \in [D]$ such that $D = (d''n - \sum_{i \in [n]} d_i)/2$, we should determine a pair $I(j) = (I_1(j), I_2(j)) \in [n]^2$ such that for every $i \in [n]$ it holds that $d_i + |I_1^{-1}(i)| + |I_2^{-1}(i)| = d''$.) ■

3 The Direct Construction

We shall make use of the edge-colored variant presented in Section 2, while relying on the fact that robustly self-ordered colored multi-graphs can be efficiently transformed into robustly self-ordered (uncolored) graphs. Actually, it will be easier to present the construction as a directed edge-colored multi-graph. Hence, we first define a variant of robust self-ordering for directed edge-colored multi-graph (see Definition 3.1), then show how to construct such multi-graphs (see Section 3.1), and finally show how to transform the directed variant into an undirected one (see Section 3.2).

The construction is based on d permutations, denoted $\pi_1, \dots, \pi_d : [n] \rightarrow [n]$, and consists of the directed edge-colored multi-graph that is naturally defined by them. Specifically, for every $v \in [n]$ and $i \in [d]$, this multi-graph contains a directed edge, denoted $(v, \pi_i(v))$, that goes from vertex v to vertex $\pi_i(v)$, and is colored i .

We prove that a sufficient condition for this edge-colored directed multi-graph, denoted G_1 , to be robustly self-ordered is that a related multi-graph is an expander. Specifically, we refer to the multi-graph $G_2 = (V_2, E_2)$ that represents the actions of these permutations on pairs of vertices of G_1 ; that is, $V_2 = \{(u, v) \in [n]^2 : u \neq v\}$ and $E_2 = \{(u, v), (\pi_i(u), \pi_i(v))\} : (u, v) \in V_2 \text{ \& } i \in [d]\}$.

The foregoing requires extending the notion of robustly self-ordered (edge-colored) multi-graphs to the directed case. The extension is straightforward and is spelled-out next, for sake of good order.

Definition 3.1 (robust self-ordering of edge-colored directed multi-graphs): *Let $G = (V, E)$ be a directed multi-graph with colored edges, where $\chi : E \rightarrow \mathbb{N}$ denotes this coloring, and let E_i denote the multi-set of directed edges colored i . We say that (G, χ) is γ -robustly self-ordered if for every permutation $\mu : V \rightarrow V$ it holds that*

$$\sum_{i \in \mathbb{N}} \left| E_i \Delta \{(\mu(u), \mu(v)) : (u, v) \in E_i\} \right| \geq \gamma \cdot |\{i \in V : \mu(i) \neq i\}|, \quad (3)$$

where $A \Delta B$ denotes the symmetric difference between the multi-sets A and B (as in Definition 2.1).

(The only difference between Definition 3.1 and Definition 2.1 is that Eq. (3) refers to the directed edges of the directed multi-graph, whereas Eq. (2) refers to the undirected edges of the undirected multi-graph.)

In Section 3.1 we present a construction of a directed edge-colored $O(1)$ -regular multi-graph that is $\Omega(1)$ -robustly self-ordered. We shall actually present a sufficient condition and a specific instantiation that satisfies it. In Section 3.2 we show how to transform any directed edge-colored multi-graph into an undirected one while preserving all relevant features; that is, bounded robustness, bounded degree, regularity, expansion, and local computability.

3.1 A sufficient condition for robust self-ordering of directed colored graphs

For any d permutations, $\pi_1, \dots, \pi_d : [n] \rightarrow [n]$, we consider two multi-graphs.

1. The **primary multi-graph** (of π_1, \dots, π_d) is a *directed* multi-graph, denoted $G_1 = ([n], E_1)$, such that $E_1 = \{(v, \pi_i(v)) : v \in [n] \ \& \ i \in [d]\}$. This directed multi-graph is coupled with an edge-coloring in which the directed edge from v to $\pi_i(v)$ is colored i .
2. The **secondary multi-graph** (of π_1, \dots, π_d) is an undirected multi-graph, denoted $G_2 = (V_2, E_2)$, such that $V_2 = \{(u, v) \in [n]^2 : u \neq v\}$ and $E_2 = \{(u, v), (\pi_i(u), \pi_i(v))\} : (u, v) \in V_2 \ \& \ i \in [d]\}$.

Recalling that we wish the secondary multi-graph to be an expander, we mention that an archetypical case is when each of the foregoing multi-graphs is a *Schreier graph* that correspond to the action of the permutation π_1, \dots, π_d on the corresponding vertex sets (i.e., $[n]$ and V_2 , respectively). See Proposition 3.3 and a wider perspective at the (paragraph at the) end of this subsection.

We now state the main result of this section, which asserts that the primary multi-graph G_1 is robustly self-ordered if the secondary multi-graph G_2 is an expander. We use the combinatorial definition of expansion: *A multi-graph $G = (V, E)$ is γ -expanding if, for every subset S of size at most $|V|/2$, there are at least $\gamma \cdot |S|$ vertices in $V \setminus S$ that neighbor some vertex in S .*

Theorem 3.2 (expansion of G_2 implies robust self-ordering of G_1): *For any $d \geq 2$ permutations, $\pi_1, \dots, \pi_d : [n] \rightarrow [n]$, if the secondary multi-graph G_2 of π_1, \dots, π_d is γ -expanding, then the primary directed multi-graph G_1 of π_1, \dots, π_d coupled with the foregoing edge-coloring is $(\gamma/2)$ -robustly self-ordered. Furthermore, G_1 (or rather the undirected multi-graph underlying G_1) is $\min(0.25, \gamma/3)$ -expanding.*

Proof: Let $\mu : [n] \rightarrow [n]$ be an arbitrary permutation, and let $T = \{v \in [n] : \mu(v) \neq v\}$ be its set of non-fixed-points. Then, the size of the symmetric difference between G_1 and $\mu(G_1)$ equals $2 \cdot \sum_{i \in [d]} |D_i|$ such that $v \in D_i$ if $(\mu(v), \mu(\pi_i(v)))$ is either not an edge in G_1 or is not colored i in it, whereas $(v, \pi_i(v))$ is an edge colored i in G_1 . Note that if $(\mu(v), \mu(\pi_i(v)))$ is not an i -colored edge in G_1 , then $\pi_i(\mu(v)) \neq \mu(\pi_i(v))$. Hence, $D_i = \{v \in [n] : \mu(\pi_i(v)) \neq \pi_i(\mu(v))\}$.

The key observation (proved next) is that *if $v \in T \setminus D_i$, then $(\pi_i(v), \pi_i(\mu(v))) \in T_2$, where $T_2 = \{(v, \mu(v)) : v \in T\}$ represents the set of replacements performed by μ .* This fact implies that if $\sum_{i \in [d]} |D_i|$ is small in comparison to $|T|$, then the set T_2 (which is a set of vertices in G_2) does not expand much, in contradiction to the hypothesis. Details follow.

Observation 3.2.1 (key observation): *For T , D_i and T_2 as defined above, if $v \in T \setminus D_i$, then $(\pi_i(v), \pi_i(\mu(v))) \in T_2$.*

Recall that $v \in T$ implies $(v, \mu(v)) \in T_2$. Observation 3.2.1 asserts that if (in addition to $v \in T$) it holds that $v \notin D_i$, then $(\pi_i(v), \pi_i(\mu(v)))$ is also in T_2 . This means that the vertices in $\{(\pi_i(v), \pi_i(\mu(v))) : v \in T \setminus D_i\}$ do not contribute to the expansion of the set T_2 in G_2 .

Proof: Since $v \notin D_i$ we have $\pi_i(\mu(v)) = \mu(\pi_i(v))$, and $\mu(\pi_i(v)) \neq \pi_i(v)$ follows, because otherwise $\pi_i(\mu(v)) = \pi_i(v)$, which implies $\mu(v) = v$ in contradiction to $v \in T$. However, $\mu(\pi_i(v)) \neq \pi_i(v)$ means that $\pi_i(v) \in T$, and $(\pi_i(v), \mu(\pi_i(v))) \in T_2$ follows. Using $\mu(\pi_i(v)) = \pi_i(\mu(v))$ again, we get $(\pi_i(v), \pi_i(\mu(v))) \in T_2$. ■

Establishing the main claim (i.e., robustness of G_1). Recall that Observation 3.2.1 implies that $\{(\pi_i(v), \pi_i(\mu(v))) : v \in T \setminus D_i\} \subseteq T_2$. On the other hand, assuming that the sequence of π_i 's contains

its own inverses (i.e., $\forall i \in [d] \exists j \in [d]$ such that $\pi_i^{-1} = \pi_j$), we observe that $\bigcup_{i \in [d]} \{(\pi_i(v), \pi_i(\mu(v))) : v \in T\}$ is the neighborhood of T_2 in the multi-graph G_2 (since $\{(\pi_i(v), \pi_i(\mu(v))) : i \in [d]\}$ is the neighbor-set of $(v, \mu(v))$ in G_2). Using the γ -expansion of the set T_2 in G_2 (while relying on $|T_2| \leq n < |V_2|/2$), it follows that

$$\sum_{i \in [d]} |\{(\pi_i(v), \pi_i(\mu(v))) : v \in T \cap D_i\}| \geq \gamma \cdot |T_2|.$$

Hence, $\sum_{i \in [d]} |D_i| \geq \gamma \cdot |T|$, and the main claim follows in this case. We reduce the general case to this special case by augmenting the sequence of π_i 's by their inverses (i.e., we add the permutations $\pi_1^{-1}, \dots, \pi_d^{-1}$, which are associated colors $d+1, \dots, 2d$). Observing that the corresponding primary graph is γ -robustly self-ordered and that it is twice more robust than the original G_1 , the claim follows.

Establishing the furthermore claim (i.e., expansion of G_1). The expansion of G_1 is shown by relating sets of vertices of G_1 to the corresponding sets of pairs in G_2 . Specifically, for and $S \subset [n]$ of size at most $n/2$, we consider the set $T = \{(u, v) \in V_2 : u, v \in S\}$, which has size $|S| \cdot (|S| - 1) \leq \frac{n}{2} \cdot (\frac{n}{2} - 1) < \frac{|V_2|}{2}$. Letting T' denote the set of neighbors of T in G_2 , and S' denote the set of neighbors of S in G_1 , on the one hand we have $|T' \setminus T| \geq \gamma \cdot |T|$ (by expansion of T in G_2), and on the other hand $|T' \setminus T| \leq 2 \cdot |S| \cdot |S' \setminus S| + |S' \setminus S| \cdot (|S' \setminus S| - 1)$. This implies $|S' \setminus S| \geq (\gamma/3) \cdot |S|$ (unless $|S| \leq 4$, which can be handled by using $|S' \setminus S| \geq 1$). ■

Primary and secondary multi-graphs based on $\text{SL}_2(p)$. Recall that $\text{SL}_2(p)$ is the multiplicative group of 2-by-2 matrices over $\text{GF}(p)$ that have determinant 1. There are several different explicit constructions of constant-size expanding generating sets for $\text{SL}_2(p)$, namely making the associated Cayley graph an expander (see, e.g., [29], [28, Thm. 4.4.2(i)], and [6]). We use any such generating set to define a directed (edge-colored) multi-graph G_1 on $p+1$ vertices, and show that the associated multi-graph on pairs, G_2 , is an expander.

Proposition 3.3 (expanding generators for $\text{SL}_2(p)$ yield an expanding secondary multi-graph): *For any prime $p > 2$, let $V = \{(1, i)^\top : i \in \text{GF}(p)\} \cup \{(0, 1)^\top\}$, and $M_1, \dots, M_d \in \text{SL}_2(p)$. For every $i \in [d]$, define $\pi_i : V \rightarrow V$ such that $\pi_i(u) = v$ if $v \in V$ is a non-zero multiple of $M_i u$. Then:*

1. Each π_i is a bijection.
2. If the Cayley multi-graph $\mathcal{C} = \mathcal{C}(\text{SL}_2(p), \{M_1, \dots, M_d\}) = (\text{SL}_2(p), \{\{M, M_i M\} : M \in \text{SL}_2(p) \& i \in [d]\})$ is an expander, then the (Schreier) multi-graph G_2 with vertex-set $P = \{(v, v') : v \in V \& v' \in V \setminus \{v\}\}$ and edge-set $\{\{(v, v'), (\pi_i(v), \pi_i(v'))\} : (v, v') \in P\}$ is an expander.

Part 1 implies that these permutations yield a primary (directed edge-colored) multi-graph on the vertex-set V , whereas Part 2 asserts that the corresponding secondary graph is an expander (if the corresponding Cayley graph is expanding). Note that $|V| = p+1$ and $|P| = (p+1)p$, whereas $|\text{SL}_2(p)| = p^3 - p = (p-1) \cdot |P|$.

Proof: Part 1 follows by observing that for every $M \in \text{SL}_2(p)$ and every vector $v \in \text{GF}(p)^2$ and scalar $\alpha \in \text{GF}(p)$ it holds that $M\alpha v = \alpha Mv$. Consequently, if for some non-zero $\alpha, \alpha' \in \text{GF}(p)$ it holds that $\alpha Mv = \alpha' Mv'$, then $Mv = M\alpha'' v'$ for $\alpha'' = \alpha'/\alpha$, which implies $v = \alpha'' v'$ (since M is invertible). Hence, $\pi_i(v) = \pi_i(v')$, for $v, v' \in V$, implies that $M_i v$ and $M_i v'$ are non-zero multiples

of the same $w \in V$, which implies $v = v'$ (since V contains a single non-zero multiple of each vector in it).

Part 2 follows by observing that the vertices of G_2 correspond to equivalence classes of the vertices of \mathcal{C} that are preserved by $\text{SL}_2(p)$, where $A, B \in \text{SL}_2(p)$ are equivalent if the columns of A are non-zero multiples of the corresponding columns of B . That is, we consider an equivalence relation, denoted \equiv , such that for $A = [A_1|A_2]$ and $B = [B_1|B_2]$ in $\text{SL}_2(p)$ it holds that $A \equiv B$ if $A_i = \alpha_i B_i$ for both $i \in \{1, 2\}$, where $\alpha_1, \alpha_2 \in [p-1]$ (and, in fact, $\alpha_2 = 1/\alpha_1$).¹⁵ By saying that these *equivalence classes are preserved by $\text{SL}_2(p)$* , we mean that, for every $A, B, M \in \text{SL}_2(p)$, if $A \equiv B$, then $MA \equiv MB$. Hence, the (combinatorial) expansion of G_2 follows from the expansion of \mathcal{C} , because the neighbors of a vertex-set $S \subseteq P$ in G_2 are the vertices of G_2 that are equivalent to T' such that T' is the set of vertices of \mathcal{C} that neighbor (in \mathcal{C}) vertices that are equivalent to vertices in S .¹⁶ ■

A simple construction. Combining Theorem 3.2 with Proposition 3.3, while using a simple pair of expanding generators (which does not yield a Ramanujan graph), we get

Corollary 3.4 (a simple robustly self-ordered primary multi-graph): *For any prime $p > 2$, let $V = \{(1, i)^\top : i \in \text{GF}(p)\} \cup \{(0, 1)^\top\}$, and consider the matrices*

$$M_1 \stackrel{\text{def}}{=} \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \quad \text{and} \quad M_2 \stackrel{\text{def}}{=} \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} \quad (4)$$

Then, for π_1 and π_2 defined as in Proposition 3.3, the corresponding primary (directed edge-colored) multi-graph is robustly self-ordered.

This follows from the fact that the corresponding Cayley graph $\mathcal{C}(\text{SL}_2(p), \{M_1, M_2\})$ is an expander [28, Thm. 4.4.2(i)].

Perspective. The foregoing construction using the group $\text{SL}_2(p)$ is a special case of a much more general family of constructions, and the elements of the proof of Proposition 3.3 follow an established theory (explained, e.g., in [26, Sec. 11.1.2]), which we briefly describe.

Let H be any finite group, and S an expanding generating set of H (i.e., the Cayley graph $\mathcal{C}(H, S)$ is an expander). Assume that H acts on a finite set V (i.e., each $h \in H$ is associated with a permutation on V , and $h'h(v) = h'(h(v))$ for every $h, h' \in H$ and $v \in V$). Then, the primary (directed edge-colored) multi-graph G_1 on vertices V can be constructed from the permutations defined by members of S . The secondary multi-graph G_2 is naturally defined by the action of S on pairs of elements in V . Finally, the expansion of $\mathcal{C}(H, S)$ implies that every connected component of G_2 is an expander.¹⁷ Thus, whenever this (Schreier) graph G_2 is connected (as it is in Proposition 3.3), one may conclude that G_1 is a directed edge-colored robustly self-ordered multi-graph.

¹⁵Recall that $\det(A) = 1 = \det(B)$, whereas $\det([\alpha_1 B_1 | \alpha_2 B_2]) = \alpha_1 \alpha_2 \cdot \det(B)$. Note that each equivalence class contains a single element of P .

¹⁶Specifically, let S have density at most half in P , and let T be the set of vertices of \mathcal{C} that are equivalent to S . Note that $|T| = (p-1) \cdot |S|$, since each equivalence class contains a single element of P . By the foregoing, the set of neighbors of T in \mathcal{C} , denoted T' , is a collection of equivalence classes of vertices of G_2 , and $|T' \setminus T| = \Omega(|T|)$ by the expansion of \mathcal{C} . It follows that the set of neighbors of S in G_2 , denoted S' , is the set of vertices that are equivalent to T' , which implies that $|S' \setminus S| = \frac{|T' \setminus T|}{p-1} = \frac{\Omega(|T|)}{p-1} = \Omega(|S|)$.

¹⁷Indeed, this was easy to demonstrate directly in the case of Proposition 3.3.

3.2 From the directed variant to the undirected one

In this section we show how to transform directed (edge-colored) multi-graphs, of the type constructed in Section 3.1, into undirected ones, while preserving all relevant features (i.e., bounded robustness, bounded degree, regularity, expansion, and local computability). The transformation is extremely simple and natural: We replace the directed edge (u, v) colored j by a 2-path with a designated auxiliary vertex $a_{u,v,j}$, while coloring the edge $\{u, a_{u,v,j}\}$ by $2j - 1$ and the edge $\{a_{u,v,j}, v\}$ by $2j$. Evidently, this colored 2-path encodes the direction of the original edge (as well as the original color).

Note that the foregoing transformation works well provided that there are no parallel edges that are colored with the same color, a condition which is satisfied by the construction presented in Section 3.1. Furthermore, since the latter construction has no vertices of (in+out) degree less than $2d \geq 4$, there is no need to mark the original vertices by self-loops. Hence, a preliminary step akin to Construction 2.2 is unnecessary here, although it can be performed in general.

Proposition 3.5 (from directed robust self-ordering to undirected robust self-ordering): *For constants $d \geq 3$ and $c \in \mathbb{N}$, let $G = (V, E)$ be a directed multi-graph in which each vertex has between three and d incident edges (in both directions), and suppose that G is coupled with an edge-coloring function $\chi: E \rightarrow [c]$ such that no parallel edges (in the same direction) are assigned the same color. Letting $E_i = \{e \in E : \chi(e) = i\}$ denote the set of directed edges colored i in G , consider the undirected multi-graph $G' = (V', E')$ such that $V' = V \cup \{a_{u,v,i} : (u, v) \in E_i\}$ and $E' = \bigcup_{j \in [2c]} E'_j$ where*

$$\begin{aligned} E'_{2i-1} &= \{\{u, a_{u,v,i}\} : (u, v) \in E_i\}, \\ E'_{2i} &= \{\{a_{u,v,i}, v\} : (u, v) \in E_i\}, \end{aligned}$$

and the edge-coloring function $\chi': E' \rightarrow [2c]$ that assigns the edges of E'_j the color j (i.e., $\chi'(e) = j$ for every $e \in E'_j$). Then, if (G, χ) is γ -robustly self-ordered (in the sense of Definition 3.1), then (G', χ') is $(\gamma/2d)$ -robustly self-ordered (in the sense of Definition 2.1), provided $\gamma \leq 1$.

We comment that the transformation of (G, χ) to (G', χ') preserves bounded robustness, bounded degree, regularity, expansion, and local computability (cf. Theorem 2.8).

Proof: The proof is analogous to the proof of Theorem 2.4, but it is much simpler because the gadgets used in the current transformation (i.e., the auxiliary vertices $a_{u,v,i}$) are much simpler.

Considering an arbitrary permutation $\mu' : V' \rightarrow V'$, we lower-bound the distance between G' and $\mu'(G')$ as a function of the number of non-fixed-points under μ' . We do so by considering the contribution of each non-fixed-point to the distance between G' and $\mu'(G')$. We first recall the fact that the vertices of V (resp., the auxiliary vertices) are uniquely identified in $\mu'(G')$ by virtue of their degree, since each vertex of V has degree at least three (in G') whereas the auxiliary vertices have degree 2.

Case 1: *Auxiliary vertices of the form $a_{u,v,i}$ that are not mapped by μ' to auxiliary vertices of the form $a_{u',v',i}$; that is, $\mu'(a_{u,v,i}) \in (V \cup \bigcup_{j \neq i} \{a_{u',v',j} : (u', v') \in E\})$.*

Each such vertex $a_{u,v,i}$ contributes at least one unit to the difference between G' and $\mu'(G')$, since the two edges incident at $a_{u,v,i}$ (in G') are colored $2i - 1$ and $2i$ respectively, whereas $\mu(a_{u,v,i})$ has either more than two edges (in G') or its two edges are colored $2j - 1$ and $2j$,

respectively, where for $j \neq i$. Hence, if the current case contains n_1 vertices, then their contribution to the distance between G' and $\mu'(G')$ is at least n_1 .

Ditto for vertices of V that are mapped by μ' to an auxiliary vertex.

Case 2: Vertices $v \in V$ such that $\mu'(v) \in V \setminus \{v\}$.

Intuitively, the hypothesis that the edge-colored directed G is robustly self-ordered, implies that such vertices contribute proportionally to the difference between the colored versions of the directed multi-graphs G and $\mu(G)$, where μ is the restriction of μ' to V . Indeed, we first assume, for simplicity, that $\mu'(V) = V$, an assumption we shall have to dispose of later. In this case, the number of tuples $((u, w), i)$ such that (u, w) is colored i in exactly one of these multi-graph (i.e., either in G or in $\mu(G)$ but not in both) is at least $\gamma \cdot |\{v \in V : \mu(v) \neq v\}|$. Assume, without loss of generality that $(u, w) \in E_i$ but either $(\mu^{-1}(u), \mu^{-1}(w)) \notin E$ (which includes the case that $\mu^{-1}(u), \mu^{-1}(w) \in V$ does not hold) or $(\mu^{-1}(u), \mu^{-1}(w)) \in E_j$ for $j \neq i$. We consider two sub-cases:

1. μ' maps some $a_{u',w',i}$ to $a_{u,w,i}$, but either u or w is not connected in G' to this $a_{u',w',i}$ via an edge with the relevant color (i.e., either $2i - 1$ or $2i$). In this sub-case we get a contribution of at least one unit, since u and w are connected to $a_{u,w,i}$ in G' .
2. A vertex not in $\{a_{u',w',i} : (u', w') \in E_i\}$ is mapped by μ' to $a_{u,w,i}$, which means that Case 1 applies for each such a tuple.

Hence, if the number of vertices in the current case is n_2 , then the number of tuples (handled by the two sub-cases) is at least $\gamma \cdot n_2$, and we get a contribution of at least $\gamma \cdot n_2$.

The foregoing description is based on the assumption that $\mu(V) = V$. If this does not hold, then we redefine μ such that $\mu(v) \notin V$ is modified such that $\mu(v) = r$ if $r \in V$ has no preimage under μ' . (Of course, each such r is only used once.) Indeed, the modified μ may be ficticiously charged with d edges per each modification, but each such modification arises due to $v \in V$ that contributes at least one unit in Case 1. Hence, the amortized over-counting of $d \cdot \gamma$ units is partially compensated by the unit contributed in Case 1.

Case 3: Auxiliary vertices of the form $a_{u,v,i}$ that are mapped by μ' to auxiliary vertices of the form $a_{u',v',i}$ for $(u',v') \neq (u,v)$; that is, $\mu'(a_{u,v,i}) \in \{a_{u',v',i} : (u',v') \in E_i \setminus \{(u,v)\}\}$.

For u, v, u', v' and i as above, if $\mu'(u) = u'$ and $\mu'(v) = v'$, then an auxiliary vertex that connects u and v in G' is mapped to an auxiliary vertex that does not connects them in $\mu'(G')$ (but rather connects the vertices u' and v' , whereas either $u' \neq u$ or $v' \neq v$). So we get a contribution of at least one unit to the difference between G' and $\mu'(G')$ (i.e., the edge incident at either u or v). Hence, the contribution is proportional to the number of non-fixed-points of the current type. Otherwise (i.e., $(\mu'(u), \mu'(v)) \neq (u', v')$), we get a vertex as in either Case 1 or Case 2, and get a proportional contribution again.

Hence, the contribution of each of these cases to the difference between G' and $\mu'(G')$ is proportional to the number of vertices involved. Specifically, if there are n_i vertices in Case i , then we get a contribution-count of at least $\gamma \cdot \sum_{i \in [3]} n_i$, where some of these contributions were possibly counted twice. The claim follows. ■

4 The Three-Step Construction

In this section we present a different construction of bounded-degree graphs that are robustly self-ordered. It uses totally different techniques than the ones utilized in the construction presented in Section 3. Furthermore, the current construction offers the flexibility of obtaining either graphs that have small connected components (i.e., of logarithmic size) or graphs that are highly connected (i.e., are expanders). Actually, one can obtain anything in-between (i.e., n -vertex graphs that consist of $s(n)$ -sized connected components that are each an expander, for any $s(n) = \Omega((\log n)/\log \log n)$). We mention that robustly self-ordered bounded-degree graphs with small connected components are used in the proof of Theorem 5.2.

As stated in Section 1.1.2, the current construction proceeds in three steps. First, in Section 4.1, we prove the existence of robustly self-ordered bounded-degree graphs, and observe that such ℓ -vertex graphs can actually be found in $\text{poly}(\ell!)$ -time (i.e., $\exp(\tilde{O}(\ell))$ -time). Next, setting $\ell = \Omega((\log n)/\log \log n)$, we use these graphs as part of 2ℓ -vertex connected components in an n -vertex (robustly self-ordered bounded-degree) graph that is constructed in $\text{poly}(n)$ -time (see Section 4.2). Lastly, in Section 4.3, we repeat this strategy using the graphs constructed in Section 4.2, and obtain exponentially larger graphs that are locally constructible.

In addition, in Section 4.4, we show that the foregoing graphs can be locally self-ordered. That is, given a vertex v in any graph $G' = (V', E')$ that is isomorphic to the foregoing n -vertex graph and oracle access to the incidence function of G' , we can find in $\text{poly}(\log n)$ -time the vertex to which this unique isomorphism maps v .

4.1 Existence

As stated above, we start with establishing the mere existence of bounded-degree graphs that are robustly self-ordered.

Theorem 4.1 (robustly self-ordered graphs exist): *For any sufficiently large constant d , there exists a family $\{G_n\}_{n \in \mathbb{N}}$ of robustly self-ordered d -regular graphs. Furthermore, these graphs are expanders.*

Actually, it turns out that random d -regular graphs are robustly self-ordered; see Theorem 6.1. Either way, given the existence of such n -vertex graphs, they can actually be found in $\text{poly}(n!)$ -time, by an exhaustive search. Specifically, for each of the possible $o((dn)!)$ graphs, we check the robust self-ordering condition by checking all $n! - 1$ relevant permutation. (The expansion condition can be checked similarly, by trying all $(0.5 + o(1)) \cdot 2^n$ relevant subsets of $[n]$.)

The proof of Theorem 4.1 utilizes a simpler probabilistic argument than the one used in the proof of Theorem 6.1. This argument (captured by Claim 4.1.1) refers to the auxiliary model of edge-colored multi-graphs (see Definition 2.1) and is combined with a transformation of this model to the original model of uncolored graphs (provided in Construction 2.3 and analyzed in Theorem 2.4). Indeed, the relative simplicity of Claim 4.1.1 is mainly due to using the edge-colored model (see digest at the end of Section 6).

Proof: To facilitate the proof, we present the construction while referring to the edge-colored model presented in Section 2. We shall then apply Theorem 2.4 and obtain a result for the original model (of uncolored simple graphs).

For $m = n/O(1)$, we shall consider $2m$ -vertex multi-graphs that consists of two m -vertex cycles, using a different color for the edges of each cycle, that are connected by $d' = O(1)$ random perfect matching, which are also each assigned a different color. (Hence, we use $2 + d'$ colors in total.) We shall show that (w.h.p.) a random multi-graph constructed in this way is robustly self-ordered (in the colored sense). (Note that parallel edges, if they exist, will be assigned different colors.) Specifically, we consider a generic $2m$ -vertex multi-graph that is determined by d' perfect matchings of $[m]$ with $\{m+1, \dots, 2m\}$. Denoting this sequence of perfect matchings by $\overline{M} = (M_1, \dots, M_{d'})$, we consider the (edge-colored) multi-graph $G_{\overline{M}}([2m], E_{\overline{M}})$ given by

$$E_{\overline{M}} = C_1 \cup C_2 \cup \bigcup_{j \in [d']} M_j$$

where $C_1 = \{\{i, i+1\} : i \in [m-1]\} \cup \{\{m, 1\}\}$
and $C_2 = \{\{m+i, m+i+1\} : i \in [m-1]\} \cup \{\{2m, m+1\}\}$

and a coloring χ in which the edges of C_j are colored j and the edges of M_j are colored $j+2$. (That is, for $i \in \{1, 2\}$, the set C_i forms a cycle of the form $((i-1)m+1, (i-1)m+2, \dots, (i-1)m+m, (i-1)m+1)$ and its edges are colored i .) Note that the $d' + 1$ edges incident at each vertex are assigned $d' + 1$ different colors.

Claim 4.1.1 (w.h.p., $G_{\overline{M}}$ is robustly self-ordered): *For some constant $\gamma > 0$, with high probability over the choice of \overline{M} , the edge-colored multi-graph $G_{\overline{M}}$ is γ -robustly self-ordered. Furthermore, it is also an expander.*

Proof: Consider an arbitrary permutation $\mu : [2m] \rightarrow [2m]$, and let $t = |\{i \in [2m] : \mu(i) \neq i\}|$. We shall show that, with probability $1 - \exp(-\Omega(dt \log m))$ over the choice of \overline{M} , the difference between the colored versions of $G_{\overline{M}}$ and $\mu(G_{\overline{M}})$ is $\Omega(t)$. Towards this end, we consider two cases.

Case 1: $|\{i \in [m] : \mu(i) \notin [m]\}| > t/4$. Equivalently, $|\{i \in [2m] : \lceil \mu(i)/m \rceil \neq \lceil i/m \rceil\}| > t/2$.

The vertices in the set $\{i \in [m] : \mu(i) \notin [m]\}$ are mapped from the first cycle to the second cycle, and so rather than having two incident edges that are colored 1 they have two incident edges colored 2. Hence, each such vertex contributes two units to the difference (between the colored versions of $G_{\overline{M}}$ and $\mu(G_{\overline{M}})$), and the total contribution is greater than $2 \cdot (t/4) \cdot 2$, where the first factor of 2 accounts also for vertices that are mapped from C_2 to C_1 .

Case 2: $|\{i \in [m] : \mu(i) \notin [m]\}| \leq t/4$. Equivalently, $|\{i \in [2m] : \lceil \mu(i)/m \rceil \neq \lceil i/m \rceil\}| \leq t/2$.

We focus on the non-fixed-points of μ that stay on their original cycle (i.e., those not considered in Case 1). Let $A \stackrel{\text{def}}{=} \{i \in [m] : \mu(i) \neq i \wedge \mu(i) \in [m]\}$ and $B \stackrel{\text{def}}{=} \{i \in \{m+1, \dots, 2m\} : \mu(i) \neq i \wedge \mu(i) \in \{m+1, \dots, 2m\}\}$. By the case hypothesis, $|A| + |B| \geq t/2$, and we may assume (without loss of generality) that $|A| \geq t/4$. As a warm-up, we first show that *each element of A contributes a non-zero number of units to the difference* (between the colored versions of $G_{\overline{M}}$ and $\mu(G_{\overline{M}})$) *with probability $1 - O(1/m)^{d'}$, over the choice of \overline{M} .*

To see this, let $\pi_j : [m] \rightarrow \{m+1, \dots, 2m\}$ be the mapping used in the j^{th} matching; that is, $M_j = \{\{i, \pi_j(i)\} : i \in [m]\}$, which means that $\pi_j(i)$ is the j^{th} match of i in $G_{\overline{M}}$ (i.e., the vertex matched to i by M_j). Then, we consider the event that *for some $j \in [d']$, the j^{th} match of $i \in [m]$ in $\mu(G_{\overline{M}})$ is different from the j^{th} match of i in $G_{\overline{M}}$* , and note that when this event

occurs i contributes to the difference (between the colored versions of $G_{\overline{M}}$ and $\mu(G_{\overline{M}})$). Note that x is the j^{th} match of i in $\mu(G_{\overline{M}})$ if and only if $\mu^{-1}(x)$ is the j^{th} match of $\mu^{-1}(i)$ in $G_{\overline{M}}$, which holds if and only if $\mu^{-1}(x) = \pi_j(\mu^{-1}(i))$ (equiv., $x = \mu(\pi_j(\mu^{-1}(i)))$). Hence, $i \in [m]$ contributes to the difference if and only if for some j it holds that $\pi_j(i) \neq \mu(\pi_j(\mu^{-1}(i)))$, because $\pi_j(i) \neq \mu(\pi_j(\mu^{-1}(i)))$ means that the edge $\{i, \pi_j(i)\}$ is colored $j+2$ in $G_{\overline{M}}$ but is not colored $j+2$ in $\mu(G_{\overline{M}})$ (since a different edge incident at i in $\mu(G_{\overline{M}})$ is colored $j+2$). Letting $\overline{\pi} = (\pi_1, \dots, \pi_{d'})$, the probability of the complementary event (i.e., i does not contribute to the difference) is given by

$$\begin{aligned} \Pr_{\overline{\pi}} [(\forall j \in [d']) \pi_j(i) = \mu(\pi_j(\mu^{-1}(i)))] &= \prod_{j \in [d']} \Pr_{\pi_j} [\pi_j(i) = \mu(\pi_j(\mu^{-1}(i)))] \\ &\leq (m-1)^{-d'}, \end{aligned}$$

where the inequality uses the hypothesis that $\mu(i) \neq i$ and $i, \mu(i) \in [m]$; specifically, fixing the value of $\pi_j(\mu^{-1}(i))$, leaves $\pi_j(i)$ uniformly distributed in $S \stackrel{\text{def}}{=} \{m+1, \dots, 2m\} \setminus \{\pi_j(\mu^{-1}(i))\}$, which means that $\Pr_{\pi_j} [\pi_j(i) = \mu(v) \mid v = \pi_j(\mu^{-1}(i))] \leq 1/|S|$ (where equality holds if $\mu(v) \in S$).

The same argument generalises to any set $I \subseteq A$ such that $I \cap \mu(I) = \emptyset$. In such a case, letting $I = \{i_1, \dots, i_{t'}\}$, we get

$$\begin{aligned} \Pr_{\overline{\pi}} [(\forall i \in I)(\forall j \in [d']) \pi_j(i) = \mu(\pi_j(\mu^{-1}(i)))] \\ &= \prod_{k \in [t']} \prod_{j \in [d']} \Pr_{\pi_j} [\pi_j(i_k) = \mu(\pi_j(\mu^{-1}(i_k)))] \mid (\forall k' \in [k-1]) \pi_j(i_{k'}) = \mu(\pi_j(\mu^{-1}(i_{k'}))) \\ &\leq (m - 2t' + 1)^{-t'd'}, \end{aligned}$$

where the inequality uses the hypothesis that $I \cap \mu(I) = \emptyset$; specifically, for each $k \in [t']$, we use the fact that $i_k \notin \{i_1, \dots, i_{k-1}, \mu^{-1}(i_1), \dots, \mu^{-1}(i_k)\}$. Hence, fixing the values of $\pi_j(i_{k'})$ for all $k' \in [k-1]$ and the values of $\pi_j(\mu^{-1}(i_{k'}))$ for all $k' \in [k]$, and denoting these values by u_1, \dots, u_{k-1} and v_1, \dots, v_k respectively, leaves $\pi_j(i_k)$ uniformly distributed in $S \stackrel{\text{def}}{=} \{m+1, \dots, 2m\} \setminus \{u_1, \dots, u_{k-1}, v_1, \dots, v_k\}$, which means that $\Pr_{\pi_j} [\pi_j(i) = \mu(v_k) \mid \text{foregoing fixing}] \leq 1/|S|$ (where equality holds if $\mu(v_k) \in S$).

Recalling that $|A| \geq t/4$ and $t \leq 2m$, we upper-bound the probability (over the choice of \overline{M}) that A contains a $t/8$ -subset A' such that $(\forall i \in A')(\forall j \in [d']) \pi_j(i) = \mu(\pi_j(\mu^{-1}(i)))$, by taking a union bound over all possible A' and using for each such A' a subset $I \subset A'$ such that $I \cap \mu(I) = \emptyset$. (So we actually take a union bound over the I 's and derive a conclusion regarding the $t/8$ -subsets A' .) Observing that $|I| \geq |A'|/3 \geq t/24$, we conclude that, with probability at most $\binom{t}{t/24} \cdot (m/2)^{d't/24} = \exp(-\Omega(d't \log m))$ over the choice of \overline{M} , the set A contains no $t/8$ -subset A' as above. This means that, with probability at most $\exp(-\Omega(d't \log m))$, less than $t/8$ of the indices $i \in A$ contribute a non-zero number of units to the difference (between the colored versions of $G_{\overline{M}}$ and $\mu(G_{\overline{M}})$).

Hence, we have shown that, for every permutations $\mu : [2m] \rightarrow [2m]$, the probability (over the choice of \overline{M}) that the size of the symmetric difference between the colored versions of $G_{\overline{M}}$ and $\mu(G_{\overline{M}})$ is smaller than $t/8$ is $\exp(-\Omega(d't \log m))$, where t is the number of non-fixed-points of μ . Letting $\gamma = 1/8$ and taking a union bound over all (non-trivial) permutations $\mu : [2m] \rightarrow [2m]$, we

conclude that the probability, over the choice of \overline{M} , that $G_{\overline{M}}$ is not γ -robustly self-ordered is at most

$$\begin{aligned} \sum_{t \in [2m]} \binom{2m}{t} \cdot \exp(-\Omega(d't \log m)) &= \sum_{t \in [2m]} \exp(-\Omega((d' - O(1)) \cdot t \log m)) \\ &= \exp(-\Omega((d' - O(1)) \cdot \log m)), \end{aligned}$$

and the claim follows (for any sufficiently large d'), while observing that, with very high probability, these multi-graphs are expanders. ■

Back to the non-colored version. We now convert the edge-colored multi-graphs $G = G_{\overline{M}}$ that are γ -robustly self-ordered into standard graphs G' that are robustly self-ordered in the original sense. This is done by using Construction 2.3 (while relying on Theorem 2.4). Recall that this transformation also preserves expansion. Actually, before invoking Construction 2.3, we augment the multi-graph G by adding a self-loop to each vertex, and color all these self-loops using a special color. Combining Claim 4.1.1 and Theorem 2.4, the current theorem follows. ■

4.2 Constructions

Having established the existence of bounded-degree graphs that are robustly self-ordered, we now turn to actually construct them. We shall use the fact that the proof of existence yields a construction that runs in time that is polynomial in the number of possible graphs. Specifically, for $\ell = \frac{O(\log n)}{\log \log n}$, we shall construct ℓ -vertex graphs in $\text{poly}(\ell^\ell)$ -time and use them in our construction of n -vertex graphs, while noting that $\text{poly}(\ell^\ell) = \text{poly}(n)$.

Theorem 4.2 (constructing robustly self-ordered graphs): *For any sufficiently large constant d , there exists an efficiently constructable family $\{G_n\}_{n \in \mathbb{N}}$ of robustly self-ordered graphs of maximum degree d . That is, there exists a polynomial-time algorithm that on input 1^n outputs the n -vertex graph $G_n = ([n], E_n)$. Furthermore, G_n consists of connected components of size $\frac{O(\log n)}{\log \log n} = o(\log n)$.*

Note that the connected components of G_n cannot be any smaller (than $\frac{O(\log n)}{\log \log n}$). This is the case because an asymmetric n -vertex bounded-degree graph, let alone a robustly self-ordered one, cannot have connected components of size $\frac{o(\log n)}{\log \log n}$ (because the number of t -vertex graphs of bounded-degree is $t^{O(t)}$).

Proof: The proof proceeds in two steps. We first use the existence of ℓ -vertex (d' -regular) expander graphs that are robustly self-ordered towards constructing a sequence of $m = \exp(\Omega(\ell \log \ell))$ bounded-degree 2ℓ -vertex graphs that are robustly self-ordered, expanding, and far from being isomorphic to one another. We construct this sequence of 2ℓ -vertex graphs in $\text{poly}(m)$ -time, using the fact that $(\ell!)^{O(1)} = \text{poly}(m)$. In the second step, we show that the $(m \cdot 2\ell)$ -vertex graph that consists of these 2ℓ -vertex graphs (as its connected components) is robustly self-ordered. Note that this graph is constructed in time that is polynomial in its size, since its size is $\Omega(m)$, whereas it is constructed in $\text{poly}(m)$ -time.¹⁸

¹⁸We mention that a slightly different construction can be based on the fact that random ℓ -vertex (d' -regular) graphs are robustly self-ordered expanders (see Theorem 6.1). In this alternative construction we find a sequence of m such graphs that are pairwise far from being isomorphic to one another. As further detailed in Remark 6.2, the analysis of the alternative construction is somewhat easier than the analysis of the construction presented below, but we need the current construction for the proof of Theorem 4.5.

Given a generic n , let $\ell = \frac{O(\log n)}{\log \log n}$, which implies that $\ell^\ell = \text{poly}(n)$. By Theorem 4.1, for all sufficiently large d' , there exist ℓ -vertex d' -regular expander graphs that are robustly self-ordered (with respect to the robustness parameter c'). Furthermore, we can find such a graph, denoted G'_ℓ , in time $\text{poly}(\ell^\ell) = \text{poly}(n)$, by scanning all ℓ -vertex d' -regular graphs and checking both the expansion and the robustness (w.r.t parameter c') conditions for each of them. Actually, for $d'' = d' + 1$, we shall also find an ℓ -vertex d'' -regular expander, denoted G''_ℓ , that is robustly self-ordered.

The construction of G_n . Using G'_ℓ and G''_ℓ , we construct an n -vertex robustly self-ordered graph, denoted G_n , that consists of $n/2\ell$ connected components that are pairwise far from being isomorphic to one another. This is done by picking $m = n/2\ell$ permutations, denoted $\pi_1, \dots, \pi_m : [\ell] \rightarrow [\ell]$, that are pairwise far-apart and constructing 2ℓ -vertex graphs such that the i^{th} such graph consist of a copy of G'_ℓ and a copy of G''_ℓ that are connected by a matching as determined by the permutation π_i . Specifically, for $G'_\ell = ([\ell], E'_\ell)$ and $G''_\ell = ([\ell], E''_\ell)$, the i^{th} connected component is isomorphic to a graph with the vertex set $[2\ell]$ and the edge set

$$E'_\ell \cup \{ \{ \ell + u, \ell + v \} : \{ u, v \} \in E'_\ell \} \cup \{ \{ v, \ell + \pi_i(v) \} : v \in [\ell] \}. \quad (5)$$

(The first two sets correspond to the copies of G'_ℓ and G''_ℓ , and the third set corresponds to the matching between these copies. Note that the vertices in $[\ell]$ have degree $d' + 1$, whereas vertices in $\{ \ell + 1, \dots, 2\ell \}$ have degree $d'' + 1 \neq d' + 1$.)

To see that this construction can be carried out in $\text{poly}(n)$ -time, we need to show that the sequence of m pairwise far-apart permutations can be determined in $\text{poly}(n)$ -time, let alone that such a sequence exists. This is the case, because we can pick the permutation sequentially (one after the other) by scanning the symmetric group on $[\ell]$ and relying on the fact that for ($i < n$ and) any fixed sequence of permutations $\pi_1, \dots, \pi_{i-1} : [\ell] \rightarrow [\ell]$ it holds that a random permutation π_i is far-apart from each of the fixed $i - 1$ permutations; that is, $\Pr_{\pi_i} [|\{ v \in [\ell] : \pi_i(v) \neq \pi_j(v) \}| = \Omega(\ell)] = 1 - o(1/n)$ for every $j \in [i - 1]$.¹⁹

Towards proving that G_n is robustly self-ordered. We now prove that the resulting graph G_n , which consists of these m connected components, is c -robustly self-ordered, where c is a universal constant (which is independent of the generic n). For starters, let's verify that G_n is self-ordered. We first note that any automorphism of G_n must map the vertices of copies of G'_ℓ (resp., G''_ℓ) to vertices of copies of G'_ℓ (resp., G''_ℓ), since these are the only vertices of degree $d' + 1$ (resp., $d'' + 1$). The connectivity of these copies implies that the automorphism must map each connected component to some connected component, which determines the m connected components. The self-ordered feature of G'_ℓ and G''_ℓ determines a unique ordering on each copy, whereas the fact the permutations (i.e., π_i 's) are different imposes that each connected component is mapped to itself (i.e., the order of the connected components is preserved). Hence, the automorphism must be trivial (and it follows that G_n is self-ordered).

An analogous argument establishes the robust self-ordering of G_n , where we use the hypothesis that G'_ℓ and G''_ℓ are expanders (rather than merely connected), the choice of the π_i 's as being far-apart (rather than merely different), and the robust self-ordering of G'_ℓ and G''_ℓ (rather than their mere self-ordering). Considering an arbitrary permutation $\mu : [n] \rightarrow [n]$, these stronger features are used to establish a lower bound on the size of the symmetric difference between G_n and $\mu(G_n)$ as follows:

¹⁹Specifically, for some $\ell' = \Omega(\ell)$, we upper-bound $\Pr_{\pi} [|\{ v \in [\ell] : \pi(v) = v \}| \geq \ell - \ell']$, where $\pi : [\ell] \rightarrow [\ell]$ is a random permutation. We do so by observing that the number of permutations that have at least $\ell - \ell'$ fixed-points is at most $\binom{\ell}{\ell'} \cdot (\ell')! = \frac{\ell!}{(\ell - \ell')!}$, whereas $(\ell - \ell')! = \exp(\Omega(\ell \log \ell)) = \omega(n)$ for any ℓ' such that $\ell - \ell' = \Omega(\ell)$.

- The fact that G'_ℓ is an expander implies that if μ splits the vertices of a copy of G'_ℓ such that ℓ' vertices are mapped to copies that are different than the other $\ell - \ell' \geq \ell'$ vertices, then this contributes $\Omega(\ell')$ units to the difference between G_n and $\mu(G_n)$. Ditto for G''_ℓ , whereas mapping a copy of G'_ℓ to a copy of G''_ℓ contributes $\Omega(\ell)$ units (per the difference in the degrees).
- The robust self-ordering of G'_ℓ and G''_ℓ implies that if μ changes the index of vertices inside a component, then this yields a proportional difference between G_n and $\mu(G_n)$.
- The distance between the π_i 's (along with the aforementioned robustness) implies that if μ changes the indices of the connected components, then each such change contributes $\Omega(\ell)$ units to the difference between G_n and $\mu(G_n)$.

The actual implementation of this sketch requires a careful accounting of the various contributions. As a first step in this direction we provide a more explicit description of G_n . We denote the set of vertices of the copy of G'_ℓ (resp., G''_ℓ) in the i^{th} connected component of G_n by $F_i = \{2(i-1)\ell + j : j \in [\ell]\}$ (resp., $S_i = \{2(i-1)\ell + \ell + j : j \in [\ell]\}$). Recall that F_i and S_i are connected by the edge-set

$$\{\{2(i-1)\ell + j, 2(i-1)\ell + \ell + \pi_i(j)\} : j \in [\ell]\} \quad (6)$$

whereas the subgraph of G_n induced by F_i (resp., S_i) has the edge-set $\{\{2(i-1)\ell + u, 2(i-1)\ell + v\} : \{u, v\} \in E'_\ell\}$ (resp., $\{\{2(i-1)\ell + \ell + u, 2(i-1)\ell + \ell + v\} : \{u, v\} \in E''_\ell\}$). In addition, let $F = \bigcup_{i \in [m]} F_i$ (resp., $S = \bigcup_{i \in [m]} S_i$).

The actual proof (that G_n is robustly self-ordered). Considering an arbitrary permutation $\mu : [n] \rightarrow [n]$, we lower-bound the distance (i.e., size of the symmetric difference) between G_n and $\mu(G_n)$ as a function of the number of non-fixed-points under μ (i.e., the number of $v \in [n]$ such that $\mu(v) \neq v$). We do so by considering the (average) contribution of every non-fixed-point to the distance between G_n and $\mu(G_n)$ (i.e., number of pairs of vertices that form an edge in one graph but not in the other). We may include the same contribution in few of the following (seven) cases, but this only means that we are double-counting the contribution by a constant factor. We first consider the set of vertices that switch sides (under μ) in the connected components.

Case 1: Vertices $v \in F$ such that $\mu^{-1}(v) \in S$. Ditto for $v \in S$ such that $\mu^{-1}(v) \in F$.

Each such vertex contributes at least one unit to the distance (between G_n and $\mu(G)$) by virtue of v having degree $d' + 1$ in G_n and strictly higher degree in $\mu(G_n)$, since vertices in F have degree $d' + 1$ (in G_n) whereas vertices in S have higher degree (in G_n).²⁰

In light of Case 1, we may focus on vertices whose “type” (or “side”) is preserved by μ^{-1} . Actually, it will be more convenient to consider the set of vertices whose “type” is preserved by μ ; that is, the set $\{v \in F : \mu(v) \in F\} \cup \{v \in S : \mu(v) \in S\}$. Next, for each $i \in [m]$, we define $\mu'(i)$ to be the index of the connected component that takes the plurality of $\mu(F_i)$; that is, $\mu'(i) \stackrel{\text{def}}{=} j$ if $|\{v \in F_i : \mu(v) \in F_j\}| \geq |\{v \in F_i : \mu(v) \in F_k\}|$ for all $k \in [m]$ (breaking ties arbitrarily). Next, we consider the set of vertices that are not mapped according to the plurality vote.

²⁰Note that v neighbors u in $\mu(G_n)$ if and only if $\mu^{-1}(v)$ neighbors $\mu^{-1}(u)$ in G_n .

Case 2: Vertices $v \in F_i$ such that $\mu(v) \in F \setminus F_{\mu'(i)}$.

For starters, suppose that $|\{v \in F_i : \mu(v) \in F_{\mu'(i)}\}| \geq \ell/2$; that is, a majority of the vertices of F_i are mapped by μ to $F_{\mu'(i)}$. In this case, by the expansion of G'_ℓ , we get a contribution that is proportional to the size of the set $F'_i \stackrel{\text{def}}{=} \{v \in F_i : \mu(v) \notin F_{\mu'(i)}\}$, because (in G_n) there are $\Omega(|F'_i|)$ edges between F'_i and the rest of F_i but there are no edges between F'_i and $F_i \setminus F'_i$ in $\mu(G_n)$. In the general case, we have to be more careful since expansion is guaranteed only for sets that have size at most $\ell/2$. In such a case we use an adequate subset of F'_i . Details follow.

Let $R_i \stackrel{\text{def}}{=} \{v \in F_i : \mu(v) \in F \setminus F_{\mu'(i)}\}$. We first note that there exists a set $J \subseteq [m] \setminus \{\mu'(i)\}$ such that $F'_i \stackrel{\text{def}}{=} \bigcup_{j \in J} \{v \in F_i : \mu(v) \in F_j\}$ satisfies $|R_i|/4 \leq |F'_i| \leq \ell/2$. Recall that the subgraph of G_n induced by F_i is an expander, and consider the edges in G_n that cross the cut between F'_i and the rest of F_i . Then, this cut has $\Omega(|F'_i|)$ edges in G_n , but there are no edges between F'_i and $F_i \setminus F'_i$ in $\mu(G_n)$, because $\mu^{-1}(F'_i) \subseteq \bigcup_{j \in J} F_j$ and $\mu^{-1}(F_i \setminus F'_i) \subseteq \bigcup_{j \in [m] \setminus J} F_j$ are not connected in G_n . Hence, the total contribution of the vertices in $\{v \in F_i : \mu(v) \in F \setminus F_{\mu'(i)}\} = R_i$ to the distance (between G_n and $\mu(G)$) is $\Omega(|F'_i|)$, which is proportional to their number (i.e., it is $\Omega(|R_i|)$).

Defining $\mu''(i)$ in an analogous manner with respect to $\mu(S_i)$, we get an analogous contribution by the expander induced by S_i . Specifically, for each $i \in [m]$, we define $\mu''(i)$ to be the index of the connected component that takes the plurality of $\mu(S_i)$; that is, $\mu''(i) \stackrel{\text{def}}{=} j$ if $|\{v \in S_i : \mu(v) \in S_j\}| \geq |\{v \in S_i : \mu(v) \in S_k\}|$ for all $k \in [m]$ (breaking ties arbitrarily). Analogously to Case 2, we consider the following vertices.

Case 3: Vertices $v \in S_i$ such that $\mu(v) \in S \setminus S_{\mu''(i)}$.

Here we get a contribution of $\Omega(|\{v \in S_i : \mu(v) \in S \setminus S_{\mu''(i)}\}|)$, where the analysis is analogous to Case 2.

Recall that if $v \in F_i$ then it holds that $v = 2(i-1)\ell + j$ for some $j \in [\ell]$, and that (in G_n) vertex v has a unique neighbor in S , which is $2(i-1)\ell + \ell + \pi_i(j) \in S_i$. It will be convenient to denote this neighbor by $\phi_i(v)$; that is, for $v \in F_i$ such that $v = 2(i-1)\ell + j$, we have $\phi_i(v) = 2(i-1)\ell + \ell + \pi_i(j) \in S_i$. The next two cases refer to vertices that are mapped by μ according to the plurality vote (e.g., $v \in F_i$ is mapped to $\mu(v) \in F_{\mu'(i)}$), but their match is not mapped accordingly (i.e., $\phi_i(v) \in S_i$ is not mapped to $S_{\mu'(i)}$).

Case 4: Vertices $v \in F_i$ such that $\mu(v) \in F_{\mu'(i)}$ but $\mu(\phi_i(v)) \notin S_{\mu'(i)}$.

(Note that the condition $v \in F_i$ and $\mu(v) \in F_{\mu'(i)}$ means that vertex v is not covered in Case 2. If $\mu''(i) = \mu'(i)$, then $\mu(\phi_i(v)) \notin S_{\mu'(i)}$ means that v is covered in Case 3, since $\phi_i(v) \in S_i$. Hence, the current case is of interest only when $\mu''(i) \neq \mu'(i)$. In particular, it is of interest when referring to vertices in the i^{th} connected component of G_n that reside in the copies of G'_ℓ and G''_ℓ and are mapped according to the plurality votes of these copies, whereas these two plurality votes are inconsistent.)

We focus on the case that a vast majority of the vertices in both F_i and S_i are mapped according to the plurality votes (i.e., $\mu'(i)$ and $\mu''(i)$), since the complementary cases are covered by Cases 2 and 3, respectively. Specifically, if either $|\{v \in F_i : \mu(v) \in [n] \setminus F_{\mu'(i)}\}| > \ell/3$

or $|\{u \in S_i : \mu(u) \in [n] \setminus S_{\mu''(i)}\}| > \ell/3$, then we get a contribution of $\Omega(\ell)$ either by Cases 1&2 or by Cases 1&3. Otherwise, it follows that

$$|\{v \in F_i : \mu(v) \in F_{\mu'(i)} \wedge \mu(\phi_i(v)) \in S_{\mu''(i)}\}| \geq \ell - 2 \cdot \ell/3$$

which implies that, if $\mu'(i) \neq \mu''(i)$, then the i^{th} connected component of G_n contributes $\ell/3$ units to the difference (between G_n and $\mu(G_n)$), since v and $\phi_i(v)$ are connected in G_n , but $\mu(v) \in F_{\mu'(i)}$ and $\mu(\phi_i(v)) \in S_{\mu''(i)}$ reside in different connected components of $\mu(G_n)$. (That is, the contribution is due to vertices v of F_i that are mapped by μ to $F_{\mu'(i)}$, while the corresponding vertices $\phi_i(v)$ of S_i (which are connected to them in G_n) are mapped by μ to $S_{\mu''(i)} \subset S \setminus S_{\mu'(i)}$, whereas $F_{\mu'(i)}$ and $S_{\mu''(i)}$ are not connected in G_n , assuming $\mu'(i) \neq \mu''(i)$.)

To conclude: The contribution of the vertices of Case 4 (to the difference between G_n and $\mu(G_n)$) is proportional to the number of these vertices (where this contribution *might* have been counted already in Cases 1, 2 and 3).

Case 5: Vertices $v \in F_i$ such that $\mu(v) \notin F_{\mu''(i)}$ but $\mu(\phi_i(v)) \in S_{\mu''(i)}$.

(Equiv., vertices $v \in S_i$ such that $\mu(v) \in S_{\mu''(i)}$ but $\mu(\phi_i^{-1}(v)) \notin F_{\mu''(i)}$.)

Analogously to Case 4, the contribution of these vertices is proportional to their number. (Analogously, this augments Case 2 only in case $\mu''(i) \neq \mu'(i)$.)

In light of Cases 2–5, we may focus on indices $i \in [m]$ such that $\mu'(i) = \mu''(i)$ and on vertices in i^{th} connected component that are mapped by μ to the $\mu'(i)^{\text{th}}$ connected component (and the same "type" per Case 1). The following case refers to such vertices that do not maintain their position in this connected component.

Case 6: Vertices $v = 2(i-1)\ell + j \in F_i$ such that $\mu(v) \in F_{\mu'(i)} \setminus \{2(\mu'(i) - 1)\ell + j\}$.

Ditto for $v = 2(i-1)\ell + \ell + j \in S_i$ such that $\mu(v) \in S_{\mu''(i)} \setminus \{2(\mu''(i) - 1)\ell + \ell + j\}$.

(This case refers to vertices in F_i that are mapped to $F_{\mu'(i)}$ but do not maintain their index in the relevant copy of G'_ℓ ; indeed, $v = 2(i-1)\ell + j$ is the j^{th} vertex of F_i , but it is mapped by μ to the k^{th} vertex of $F_{\mu'(i)}$ (i.e., $\mu(v) = 2(\mu'(i) - 1)\ell + k$) such that $k \neq j$.)

Fixing i , let $C \stackrel{\text{def}}{=} \{v = 2(i-1)\ell + j \in F_i : \mu(v) \in F_{\mu'(i)} \setminus \{2(\mu'(i) - 1)\ell + j\}\}$ denote the set of vertices considered in this case, and $D = \{v \in F_i : \mu(v) \notin F_{\mu'(i)}\}$ denote the set of vertices that we are going to discount for. As a warm-up, consider first the case that $D = \emptyset$. In this case, by the robust self-ordering of G'_ℓ , the contribution of the vertices in C to the difference between G_n and $\mu(G_n)$ is $\Omega(|C|)$.

In the general case (i.e., where D may not be empty), we get a contribution of $\Omega(|C|) - d' \cdot |D|$, where the second term compensates for the fact that the vertices of D were moved outside of this copy of G'_ℓ and replaced by different vertices that may have different incidences. Letting c' be the constant hidden in the Ω -notation, we get a contribution of at least $c' \cdot |C| - d' \cdot |D|$, which is at least $c' \cdot |C|/2$ if $|D| \leq c' \cdot |C|/2d'$. On the other hand, if $|D| > c' \cdot |C|/2d'$, then we get a contribution of $\Omega(|D|) = \Omega(|C|)$ by Cases 1–2.

Hence, in both sub-cases we have a contribution of $\Omega(|C|)$ to the difference between G_n and $\mu(G_n)$.

The same analysis applies to $\{v = 2(i-1)\ell + \ell + j \in S_i : \mu(v) \in S_{\mu''(i)} \setminus \{2(\mu''(i) - 1)\ell + \ell + j\}\}$, where we use the robust self-ordering of G''_ℓ and Cases 1&3.

Lastly, we consider vertices that do not fall into any of the prior cases. Such vertices maintain their type, are mapped with the plurality vote of their connected component, which is consistent among its two parts (i.e., μ' and μ''), and maintain their position in that component. Hence, *the hypothesis that they are not fixed-points of μ can only be attributed to the fact that these vertices are mapped to a connected component with a different index.*

Case 7: Vertices $v \in F_i$ such that both $\mu(v) \in F_{\mu'(i)} \setminus F_i$ and $\mu(\phi_i(v)) \in S_{\mu''(i)} \setminus S_i$ hold.

(We may assume that $\mu'(i) \neq i$ and $\mu''(i) \neq i$, since otherwise this set is empty. We may also assume that $\mu'(i) = \mu''(i)$, since the complementary case was covered by Cases 4 and 5. Hence, we focus on pairs of vertices that are matched in the i^{th} connected component of G_n and are mapped by μ to the k^{th} component of G_n such that $k \neq i$.)

For every $i \neq k$, let $\Delta_{i,k} = \{j \in [\ell] : \pi_i(j) \neq \pi_k(j)\}$ be the sets on which π_i and π_k differ. (Note that if for every $v = 2(i-1)\ell + j \in F_i$ it holds that $\mu(v) = 2(k-1)\ell + j$ and $\mu(\phi_i(v)) = 2(k-1)\ell + \pi_i(j)$ (equiv., $\mu(2(i-1)\ell + \ell + \pi_i(j)) = 2(k-1)\ell + \pi_i(j)$), then we get a contribution of $|\Delta_{i,k}|$ to the difference between G_n and $\mu(G_n)$.)

Fixing i , let $D = D_1 \cup D_2$ such that

$$\begin{aligned} D_1 &= \{v \in F_i : \mu(v) \notin F_{\mu'(i)} \vee \mu(v + \ell) \notin S_{\mu''(i)}\} \\ D_2 &= \left\{ v = 2(i-1)\ell + j \in F_i : \begin{array}{l} \mu(v) \in F_{\mu'(i)} \setminus \{2(\mu'(i) - 1)\ell + j\} \\ \vee \mu(\phi_i(v)) \in S_{\mu''(i)} \setminus \{2(\mu''(i) - 1)\ell + \ell + \pi_i(j)\} \end{array} \right\} \end{aligned}$$

(Recall that $\phi_i(2(i-1)\ell + j) = 2(i-1)\ell + \ell + \pi_i(j)$. The set D_1 accounts for the vertices covered in Cases 2&3, whereas D_2 accounts for the vertices covered in (the two sub-cases of) Case 6.)

As a warm-up, consider first the case that $D = \emptyset$. In this case, assuming $\mu'(i) = \mu''(i) \neq i$, we get a contribution of $|\Delta_{i,\mu'(i)}| = \Omega(\ell)$ (to the difference between G_n and $\mu(G_n)$). This contribution is due to the difference in the edges that match $F_{\mu'(i)}$ and $S_{\mu'(i)}$ in G_n and the edges that match F_i and S_i in G_n , where $|\Delta_{i,\mu'(i)}| = \Omega(\ell)$ is due to the fact that the permutations (i.e., π_k 's) are far-apart. The hypothesis $D_1 = \emptyset$ means that all vertices of F_i (resp., of S_i) are mapped to $F_{\mu'(i)}$ (resp., to $S_{\mu''(i)} = S_{\mu'(i)}$), whereas $D_2 = \emptyset$ means that these vertices preserves their order within the two parts of the connected component.

The general case (i.e., where D may not be empty) requires a bit more care. Suppose that the π_k 's are γ -apart; that is, $|\Delta_{k',k}| > \gamma \cdot \ell$ for every $k' \neq k$. We focus on the case that a vast majority of the vertices in both F_i and S_i are mapped according to the plurality votes (i.e., $\mu'(i)$ and $\mu''(i)$), since the complementary cases are covered by Cases 2 and 3, respectively. Specifically, if $|D_1| > \gamma\ell/3$, then we get a contribution of $\Omega(\ell)$ by either Case 2 or Case 3. Likewise, if $|D_2| > \gamma\ell/3$, then we get a contribution of $\Omega(\ell)$ by Case 6. So, assuming $\mu'(i) = \mu''(i) \neq i$, we are left with the case that

$$|\{v = 2(i-1)\ell + j \in F_i \setminus D : j \in \Delta_{i,\mu'(i)}\}| \geq \gamma\ell - 2\gamma\ell/3.$$

In this case, we get a contribution of at least $\gamma\ell/3$ to the difference between G_n and $\mu(G_n)$. This contribution is due to the difference in the edges that match $F_{\mu'(i)}$ and $S_{\mu'(i)}$ in G_n and the edges that match F_i and S_i in G_n , where edges that have an endpoint (or its ϕ_i -mate) in D were discarded. Specifically, letting $k = \mu'(i) = \mu''(i) \neq i$, the pair $(v, w) =$

$(2(i-1)\ell + j, 2(i-1)\ell + \ell + \pi_i(j)) \in F_i \times S_i$ contributes to the difference if $j \in \Delta_{i,k}$ and both $\mu(v) = 2(k-1)\ell + j \in F_k$ and $\mu(w) = 2(k-1)\ell + \ell + \pi_i(j) \in S_k$ hold (i.e., $v \notin D_1$ and $v, \phi_i^{-1}(w) \notin D_2$).²¹ Indeed, in this case $\{v, w\}$ is an edge in G_n but $\{v, w\}$ is not an edge in $\mu^{-1}(G_n)$. (Hence, if the number of vertices of this case is $\Omega(|\{u \in [n] : \mu(u) \neq u\}|)$, then the difference between G_n and $\mu^{-1}(G_n)$ is $\Omega(|\{u \in [n] : \mu(u) \neq u\}|)$, and the same holds with respect to the difference between $\mu(G_n)$ and G_n .)

Combining all these cases, we get a total contribution that is proportional to $|\{v \in [n] : \mu(v) \neq v\}|$, where we might have counted the same contribution in several different cases. Since the number of cases is a constant, the theorem follows. \blacksquare

Digest: Using large collections of pairwise far apart permutations. The construction presented in the proof of Theorem 4.2 utilizes a collection of $(\ell!)^{\Omega(1)}$ permutations over $[\ell]$ that are pairwise far-apart (i.e., every two permutations differ on $\Omega(\ell)$ inputs). Such a collection is constructed in $\tilde{O}(\ell!)$ -time by an iterative exhaustive search, where the permutations are selected iteratively such that in each iteration we find a permutation that is far from permutations that were included in previous iterations. We mention that in Section 4.3 we shall use a collection of $\exp(\Omega(\ell))$ such permutations that is locally computable (i.e., given the index of a permutation we find its explicit description in polynomial time). We also mention that, in follow-up work [23], we provided a locally computable collection of $(\ell!)^{\Omega(1)}$ that are pairwise far-apart.

Digest: Combining two robustly self-ordered graphs. One ingredient in the proof of Theorem 4.2 is forming connected components that consist of two robustly self-ordered graphs that have different vertex degrees and are connected by a bounded-degree bipartite graph. Implicit in the proof is the fact that the resulting graph is robustly self-ordered graph.

Claim 4.3 (combining two $\Omega(1)$ -robustly self-ordered graphs): *For $i \in \{1, 2\}$ and constant $\gamma > 0$, let $G_i = (V_i, E_i)$ be an γ -robustly self-ordered graph, and consider a graph $G = (V_1 \cup V_2, E_1 \cup E_2 \cup E)$ of maximum degree d such that E contain edges with a single vertex in each V_i ; that is, G consists of G_1 and G_2 and an arbitrary bipartite graph that connects them. If the maximum degree in G of each vertex in V_1 is strictly smaller than the minimum degree of each vertex in V_2 , then G is $\gamma/(4d+1)$ -robustly self-ordered.*

Proof Sketch: For an arbitrary permutation $\mu : V \rightarrow V$, let T denote the set of its non-fixed-points, and consider the following two cases.

Case 1: More than $t = \gamma' \cdot |T|$ vertices are mapped by μ from G_1 to G_2 , where $\gamma' = \gamma/(4d+1)$.

In this case, we get a contribution of at least one unit per each such vertex, due to the difference in the degrees between V_1 and V_2 .

Case 2: at most t vertices are mapped by μ from G_1 to G_2 .

In this case, letting T_i denote the set of non-fixed vertices in G_i that are mapped by μ to G_i , we get a contribution of at least $\sum_{i=1,2} (\gamma \cdot |T_i| - d \cdot t)$ units, where the negative term is due to possible change in the incidence with vertices in $T \setminus T_i$. Hence, the total contribution in this case is at least $\gamma \cdot (|T| - 2t) - 2d \cdot t \geq \gamma \cdot |T| - 4dt = \gamma' \cdot |T|$.

²¹Recall that $\phi_i^{-1}(w) = \phi^{-1}((2(i-1)\ell + \ell + \pi_i(j))) = 2(i-1)\ell + j = v$.

The claim follows. ■

Regaining regularity and expansion. While Theorem 4.2 achieves our main objective, it is useful towards some applications (see, e.g., the proof of Theorem 4.5) to obtain this objective with graphs that are both regular and expanding. This is achieved by applying Theorem 2.6. Hence, we have.

Theorem 4.4 (Theorem 4.2, revised): *For any sufficiently large constant d , there exists an efficiently constructable family $\{G_n\}_{n \in \mathbb{N}}$ of robustly self-ordered d -regular expander graphs. That is, there exists a polynomial-time algorithm that on input 1^n outputs the n -vertex graph G_n .*

4.3 Strong (i.e., local) constructions

While Theorem 4.4 provides an efficient construction of robustly self-ordered d -regular expander graphs, we seek a stronger notion of constructability. Specifically, rather than requiring that the graph be constructed in time that is polynomial in its size, we require that the neighbors of each given vertex can be found in time that is polynomial in the vertex’s name (i.e., time that is polylogarithmic in the size of the graph). We call such graphs **locally constructable** (and comment that the term “strongly explicit” is often used in the literature).

Theorem 4.5 (locally constructing robustly self-ordered graphs): *For any sufficiently large constant d , there exists a locally constructable family $\{G_n = ([n], E_n)\}_{n \in \mathbb{N}}$ of robustly self-ordered d -regular graphs. That is, there exists a polynomial-time algorithm that on input n and $v \in [n]$ outputs the list of neighbours of vertex v in G_n . Furthermore, the graphs are either expanders or consist of connected components of logarithmic size.*

(Indeed, this establishes Theorem 1.3.) We comment that using the result of [23], we can also get connected components of sub-logarithmic size, as in Theorem 4.2.²²

Proof: We employ the idea that underlies the proof of Theorem 4.2, while starting with an *efficiently constructable family* of robustly self-ordered graphs (as provided by Theorem 4.4) rather than with the mere existence of a family of such graphs (equiv., with ℓ -vertex graphs that can be constructed in $\text{poly}(\ell)$ -time). We use a slightly larger setting of ℓ , which allows us to use a collection of $\exp(\Omega(\ell))$ pairwise-far-apart permutations (rather than a collection of $\exp(\Omega(\ell \log \ell))$ such permutations). Lastly, we apply the same transformation as in the proof of Theorem 4.4 (so to regain regularity and expansion). Details follow.

Given a generic n , let $\ell = O(\log n)$, which implies that $\exp(\ell) = \text{poly}(n)$. By Theorem 4.4, for all sufficiently large d' , we can construct ℓ -vertex d' -regular expander graphs that are robustly self-ordered (with respect to the robustness parameter c) in $\text{poly}(\ell)$ -time. Again, we shall use two such graphs: a d' -regular graph, denoted $G'_\ell = ([\ell], E'_\ell)$, and a d'' -regular graph, denoted $G''_\ell = ([\ell], E''_\ell)$, where $d'' = d' + 1$.

Using G'_ℓ and G''_ℓ , we construct an n -vertex robustly self-ordered graph, denoted G_n , that consists of $n/2\ell$ connected components that are pairwise far from being isomorphic to one another.

²²Specifically, the result of [23] provides a construction of a collection of $L = \exp(\Omega(\ell \log \ell))$ permutations over $[\ell]$ that are pairwise far-apart along with a polynomial-time algorithm that, on input $i \in [L]$, returns a description of the i^{th} permutation (i.e., the algorithm should run in $\text{poly}(\log L)$ -time). Using this algorithm, we can afford to set $\ell = \frac{O(\log n)}{\log \log n}$ as in Theorem 4.2.

This is done by picking $m = n/2\ell$ permutations, denoted $\pi_1, \dots, \pi_m : [\ell] \rightarrow [\ell]$, that are pairwise far-apart, and constructing 2ℓ -vertex graphs such that the i^{th} such graph consist of a copy of G'_ℓ and a copy of G''_ℓ that are connected by a matching as determined by the permutation π_i (as detailed in Eq. (7)).

Using the fact that $m < 2^\ell$ (rather than $m = \exp(\Theta(\ell \log \ell))$), we can construct each of these permutations in $\text{poly}(\ell)$ -time by using sequences of disjoint traspositions determined via a good error correcting code. Specifically, for $k = \log_2 m < \log_2 n$, we use an error correcting code $C : \{0, 1\}^k \rightarrow \{0, 1\}^\ell$ of constant rate (i.e., $\ell = O(k)$) and linear distance (i.e., the codewords are $\Omega(\ell)$ bits apart from each other), and let $\pi_i(2j - 1) = 2j - 1 + C(i)_j$ and $\pi_i(2j) = 2j - C(i)_j$, where $i \in [m] = [2^k] \equiv \{0, 1\}^k$ and $j \in [\ell/2]$. (That is, the i^{th} permutation switches the pair $(2j - 1, 2j) \in [\ell]^2$ if and only if the j^{th} bit in the i^{th} codeword is 1, where $C(i)$ is considered the i^{th} codeword.)

Like in the proof of Theorem 4.2, the i^{th} connected component of G_n is isomorphic to a graph with the vertex set $[2\ell]$ and the edge set

$$E'_\ell \cup \{\{\ell + u, \ell + v\} : \{u, v\} \in E''_\ell\} \cup \{\{v, \ell + \pi_i(v)\} : v \in [\ell]\}. \quad (7)$$

The key observation is that, for every $i \in [m]$ and $j \in [\ell]$, the neighborhood of the j^{th} (resp., $(\ell + j)^{\text{th}}$) vertex in the i^{th} connected component of the n -vertex graph G_n is determined by G'_ℓ and $\pi_i(j)$ (resp., by G''_ℓ and $\pi_i^{-1}(j)$), which in turn are constructible in $\text{poly}(\ell)$ -time. Hence, the neighborhood of each vertex in G_n can be found in $\text{poly}(\ell)$ -time. This implies local constructability, since $\ell = O(\log n)$.

The fact that G_n is robustly self-ordered was already established in the proof of Theorem 4.2, which is oblivious of the permutations used as long as any pair of permutations disagrees on $\Omega(\ell)$ points. Lastly, we may obtain regularity and expansion by applying Theorem 2.6. ■

4.4 Local self-ordering

Recall that, by Definition 1.1, a graph $G = ([n], E)$ is called *self-ordered* if for every graph $G' = (V', E')$ that is isomorphic to G there exists a unique bijection $\phi : V' \rightarrow [n]$ such that $\phi(G') = G$. One reason for our preferring the term “self-ordered” over the classical term “asymmetric” is that we envision being given such an isomorphic copy $G' = (V', E')$ and asked to find its unique isomorphism to G , which may be viewed as ordering the vertices of G' according to (their index in) G . The task of finding this unique isomorphism will be called *self-ordering G' according to G* or *self-ordering G'* (when G is clear from the context).

Evidently, the task of self-ordering a given graph G' according to a (self-ordered) graph G that can be efficiently constructed reduces to testing isomorphism. When the graphs have bounded-degree the latter task can be performed in polynomial-time [30]. These are general facts that do apply also to the robustly self-ordered graph G_n constructed in the proof of Theorem 4.5. However, in light of the fact that the graph G_n is *locally* constructable, we can hope for more. Specifically, it is natural to ask if we can perform self-ordering of a graph G' that is isomorphic to G_n in a *local* manner; that is, given a vertex in G' (and oracle access to the incidence function of G'), can we find the corresponding vertex in G_n in $\text{poly}(\log n)$ -time? Let us define this notion formally.

Definition 4.6 (locally self-ordering a self-ordered graph): *We say that a self-ordered graph $G = ([n], E)$ is locally self-ordered if there exists a polynomial-time algorithm that, given a vertex v in*

any graph $G' = (V', E')$ that is isomorphic to G and oracle access to the incidence function of G' , finds $\phi(v) \in [n]$ for the unique bijection $\phi : V' \rightarrow [n]$ such that $\phi(G') = G$ (i.e., the unique isomorphism of G' to G).

Indeed, the isomorphism ϕ orders the vertices of G' in accordance with the original (or target) graph G . We stress that the foregoing algorithm works in time that is polynomial in the description of a vertex (i.e., $\text{poly}(\log n)$ -time), which is polylogarithmic in the size of the graph (i.e., n). We show that such algorithms exist for the graphs constructed in the proof of Theorem 4.5.

Theorem 4.7 (locally self-ordering the graphs of Theorem 4.5): *For any sufficiently large constant d , there exists a locally constructable family $\{G_n = ([n], E_n)\}_{n \in \mathbb{N}}$ of robustly self-ordered d -regular graphs that are locally self-ordered. Furthermore, the graphs are either expanders or consist of connected components of logarithmic size.*

As in Theorem 4.5, we can obtain connected components of sub-logarithmic size by using [23].

Proof: We first consider the version that yields n -vertex graphs that consist of connected components of logarithmic size. The basic idea is that it we can afford reconstructing the connected component in which the input vertex reside, and this allows us both to determine the index of the vertex in this connected component as well as the index of the component in the graph. Specifically, on input a vertex v in a graph G' that is isomorphic to G_n , we proceed as follows.

1. Using queries to the incidence function of G' , we explore and retrieve the entire 2ℓ -vertex connected component in which v resides, where $\ell = \log_2 n$.

Recall that this connected component consists of (copies of) two ℓ -vertex regular graphs, denoted G'_ℓ and G''_ℓ , that are connected by a matching. Furthermore, these graphs have different degrees and are each (robustly) self-ordered.

2. Relying on the different degrees, we identify the foregoing partition of this 2ℓ -vertex component into two ℓ -vertex (self-ordered) graphs, denoted A_v and B_v , where A_v (resp., B_v) is isomorphic to G'_ℓ (resp., G''_ℓ).
3. Relying on the self-ordering of G'_ℓ (resp., G''_ℓ), we order the vertices of A_v (resp., B_v). This is done by constructing G'_ℓ (resp., G''_ℓ), and using an isomorphism tester. The order of the vertices in A_v and B_v also determines the permutation that defines the matching between the two graphs.
4. Relying on the correspondence between the permutations used in the construction and codewords of a good error-correcting code, we decode the relevant codeword (i.e., this is decoding without error). This yields the index of the permutation in the collection, which equals the index of the connected component.

Note that this refers to the basic construction that was presented in the proof of Theorem 4.5, before it was transformed to an expander and made regular. Recall that both transformations are performed by augmenting the graph with auxiliary edges that are assigned a different color than the original edges, and that edges with different colors are later replaced by copies of different (constant-size) gadgets. These transformations do not hinder the local self-ordering procedure described above, since it may identify the original graph (and ignore the gadgets that replace other edges). The claim follows. ■

Local reversed self-ordering. While *local self-ordering* a (self-ordered) graph seems *the natural local version* of self-ordering the graph, an alternative notion called *local reversed self-ordering* will be defined and studied next (and used in Section 5). Both notions refer to a self-ordered graph, denoted $G = ([n], E)$, and to an isomorphic copy of it, denoted $G' = (V', E')$; that is, $G = \phi(G')$ for a (unique) bijection $\phi : V' \rightarrow [n]$. While local self-ordering is the task of finding the index of a given vertex of G' according to G (i.e., given $v \in V'$, find $\phi(v) \in [n]$), local reversed self-ordering is the task of finding the vertex of G' that has a given index in G (i.e., given $i \in [n]$, find $\phi^{-1}(i) \in V'$). In both cases, the graph G is locally constructible and we are given oracle access to the incidence function of G' . In addition, in the reversed task, we assume that the algorithm is given an arbitrary vertex in G' , since otherwise there is no hope to hit any element of V' .²³

Definition 4.8 (locally reversed self-ordering): *We say that a self-ordered graph $G = ([n], E)$ is locally reversed self-ordered if there exists a polynomial-time algorithm that, given $i \in [n]$ and oracle access to the incidence function of a graph $G' = (V', E')$ that is isomorphic to G and an arbitrary vertex $s \in V'$, finds $\phi^{-1}(i) \in V'$ for the unique bijection $\phi : V' \rightarrow [n]$ such that $\phi(G') = G$ (i.e., the unique isomorphism of G' to G).*

We stress that the foregoing algorithm works in time that is polynomial in the description of a vertex (i.e., $\text{poly}(\log n)$)-time), which is polylogarithmic in the size of the graph (i.e., n). We show that such algorithms exist for variants of the graphs constructed in the proof of Theorem 4.5. In fact, we show a more general result that refers to any graph that is locally self-ordered and for which short paths can be locally found between any given pair of vertices.

Theorem 4.9 (sufficient conditions for locally reversed self-ordering of graphs): *Suppose that $\{G_n = ([n], E_n)\}_{n \in \mathbb{N}}$ is a family of bounded degree graphs that is locally self-ordered. Further suppose that given $v, u \in [n]$, one can find in polynomial-time a path from u to v in G_n . Then, $\{G_n = ([n], E_n)\}_{n \in \mathbb{N}}$ is locally reversed self-ordered.*

We mention that any family of robustly self-ordered graphs that is locally self-ordered can be transformed into one that also supports locally finding short paths. This is done by superimposing the graphs of this family with graphs that supports locally finding short paths, while using different colors for the edges of the two graphs and later replacing these colored edges by gadgets (as done in Section 2.1). We also mention that applying degree reduction to the hyper-cube (i.e., replacing the original vertices with simple cycles) yields a graph that supports locally finding short paths.²⁴

Proof: On input $i \in [n]$ and $s \in V'$, and oracle access to the incidence function of a graph $G' = (V', E')$ that is isomorphic to G_n , we proceed as follows.

1. Using the local self-ordering algorithm, we find $i_0 = \phi(s)$, where $\phi : V' \rightarrow [n]$ is the unique bijection satisfying $\phi(G') = G$.

²³Needless to say, this is not needed in case $V' = [n]$, which is the case that is used in Section 5.

²⁴For any $\ell \in \mathbb{N}$, the resulting graph consists of the vertex-set $\{\langle x, i \rangle : x \in \{0, 1\}^\ell \ \& \ i \in [\ell]\}$ and edges that connect $\langle x, i \rangle$ to $\langle x \oplus 0^{i-1}10^{\ell-i}, i \rangle$ and to $\langle x, i+1 \rangle$, where $\ell+1$ stands for 1. For simplicity of exposition, we also add self-loops on all vertices. Then, given $\langle x, i \rangle$ and $\langle y, j \rangle$, we can combine the 2ℓ -path that goes from $\langle x, i \rangle$ to $\langle y, i \rangle$ with the $|j-i|$ -path that goes from $\langle y, i \rangle$ to $\langle y, j \rangle$, where the odd steps on the first path move from $\langle z, k \rangle$ to $\langle z \oplus 0^{i-1}10^{\ell-i}, k \rangle$ (or stay in place) and the even steps (on this path) move from $\langle z, k \rangle$ to $\langle z, k+1 \rangle$.

2. Using the path-finding algorithm for G , we find a $\text{poly}(\log n)$ -long path from i_0 to i in G .
Let ℓ denote the length of the path, and denote its intermediate vertices by $i_1, \dots, i_{\ell-1}$; that is, the full path is $i_0, i_1, \dots, i_{\ell-1}, i_\ell = i$.
3. Using the local self-ordering algorithm (and our oracle access to G'), we iteratively find the corresponding vertices in G' , denoted v_1, \dots, v_ℓ , where $v_j = \phi^{-1}(i_j)$.

For $j = 1, \dots, \ell$, we find $v_j \stackrel{\text{def}}{=} \phi^{-1}(i_j)$ as follows. First, using queries to the incidence function of G' , we find all neighbors (in G') of v_{j-1} , where $v_0 \stackrel{\text{def}}{=} s$ (and, indeed, $v_0 = \phi^{-1}(i_0)$). Next, using the local self-ordering algorithm, we find the indices of all these vertices in G ; that is, for every vertex w that neighbors v_{j-1} , we find $\phi(w)$. Last, we set v_j to be the neighbor that has index i_j in G ; that is, v_j satisfies $\phi(v_j) = i_j$.

Hence, v_ℓ is the desired vertex; that is, v_ℓ satisfies $\phi(v_\ell) = i_\ell = i$.

Assuming that the local self-ordering algorithm has query complexity $q(n)$, that the paths found in G have length at most $\ell(n)$, and that d is the degree bound, the query complexity of our reversed self-ordering algorithm is $(1 + \ell(n) \cdot d) \cdot (q(n) + 1)$, where we count both our direct queries to the incidence function of G and the queries performed by the local self-ordering algorithm. Similar considerations apply to the time complexity. ■

Corollary 4.10 (a version of Theorem 4.7 supporting local reversed self-ordering): *For any sufficiently large constant d , there exists a locally constructable family $\{G_n = ([n], E_n)\}_{n \in \mathbb{N}}$ of robustly self-ordered graphs of maximum degree d that are both locally self-ordered and locally reversed self-ordered.*

The corollary follows by combining Theorem 4.7 with Theorem 4.9, while using the augmentation outlined following the statement of Theorem 4.9. We mention that Corollary 4.10 will be used in Section 5.

5 Application to Testing Bounded-Degree Graph Properties

Our interest in efficiently constructable bounded-degree graphs that are robustly self-ordered was triggered by an application to property testing. Specifically, we observed that such constructions can be used for proving a linear lower bound on the query complexity of testing an *efficiently recognizable* graph property in the *bounded-degree graph model*.

It is well known that 3-Colorability has such a lower bound [3], but this set is NP-complete. On the other hand, linear lower bounds on the query complexity of testing efficiently recognizable properties of *functions* (equiv., sequences) are well known (see [19, Sec. 10.2.3]). So the idea was to transport the latter lower bounds from the domain of functions to the domain of bounded-degree graphs, and this is where efficient constructions of robustly self-ordered bounded-degree graphs come into play. (We mention that an alternative way of obtaining the desired lower bound was outlined in [17, Sec. 1], see details below.)

More generally, the foregoing transportation demonstrates a general methodology of transporting lower bounds that refer to testing binary strings to lower bounds regarding testing graph properties in the bounded-degree graph model. The point is that strings are ordered objects, whereas graphs properties are effectively sets of unlabeled graphs, which are unordered objects. Hence, we

need to make the graphs (in the property) ordered, and furthermore make this ordering robust in the very sense that is reflected in Definition 1.2. Essentially, we provide a reduction of testing a property of strings to testing a (related) property of graphs.

We apply this methodology to obtain a subexponential separation between the complexities of testing and tolerant testing of graph properties in the bounded-degree graph model. This result is obtained by transporting an analogous result that was known for testing binary strings [15]. In addition to using a reduction from tolerantly testing a property of strings to tolerantly testing a property of graphs, this transportation also uses a reduction in the opposite direction, which relies on the local computation features asserted in Corollary 4.10.

Organization of this section. We start with a brief review of the bounded-degree graph model for testing graph properties. Next, we prove the aforementioned linear lower bound on the query complexity of testing an efficiently recognizable property, and later we abstract the reduction that underlies this proof. Observing that this reduction applies also to tolerant testing, and presenting a reduction in the opposite direction, we derive the aforementioned separation between testing and tolerant testing.

5.1 Background

Property testing refers to algorithms of sublinear query complexity for *approximate decision*; that is, given oracle access to an object, these algorithms (called testers) distinguish objects that have a predetermined property from objects that are far from the property. Different models of property testing arise from different query access and different distance measures.

In the last couple of decades, the area of property testing has attracted significant attention (see, e.g., [16]). Much of this attention was devoted to testing graph properties in a variety of models including the dense graph model [19], and the bounded-degree graph model [21] (surveyed in [16, Chap. 8] and [16, Chap. 9], resp.). In this section, we refer to the bounded-degree graph model, in which graphs are represented by their incidence function and distances are measured as the ratio of the number of differing incidences to the maximal number of edges.

Specifically, for a degree bound $d \in \mathbb{N}$, we represent a graph $G = ([n], E)$ of maximum degree d by the incidence function $g : [n] \times [d] \rightarrow [n] \cup \{0\}$ such that $g(v, i)$ indicates the i^{th} neighbor of v (where $g(v, i) = 0$ indicates that v has less than i neighbors). The distance between the graphs $G = ([n], E)$ and $G' = ([n], E')$ is defined as the size of the symmetric difference between E and E' over $dn/2$.

A tester for a property Π is given oracle access to the tested object, where here oracle access to a graph means oracle access to its incidence function. In addition, such a tester is given a size parameter n (i.e., the number of vertices in the graph), and a **proximity parameter**, denoted $\epsilon > 0$. Tolerant testers, introduced in [31] (and briefly surveyed in [16, Sec. 12.1]), are given an additional parameter, $\eta < \epsilon$, which is called the **tolerance parameter**.

Definition 5.1 (testing and tolerant testing graph properties in the bounded-degree graph model): *For a fixed degree bound d , a tester for a graph property Π is a probabilistic oracle machine that, on input parameters n and ϵ , and oracle access to an n -vertex graph $G = ([n], E)$ of maximum degree d , outputs a binary verdict that satisfies the following two conditions.*

1. *If $G \in \Pi$, then the tester accepts with probability at least $2/3$.*

2. If G is ϵ -far from Π , then the tester accepts with probability at most $1/3$, where G is ϵ -far from Π if for every n -vertex graph $G' = ([n], E') \in \Pi$ of maximum degree d it holds that the size of the symmetric difference between E and E' has cardinality that is greater than $\epsilon \cdot dn/2$.

A tolerant tester is also given a tolerance parameter η , and is required to accept with probability at least $2/3$ any graph that is η -close to Π (i.e., not η -far from Π).²⁵

We stress that a graph property is defined as a property that is preserved under isomorphism; that is, if $G = ([n], E)$ is in the graph property Π , then all its isomorphic copies are in the property (i.e., $\pi(G) \in \Pi$ for every permutation $\pi : [n] \rightarrow [n]$). The fact that we deal with graph properties (rather than with properties of functions) is the source of the difficulty (of transporting results from the domain of functions to the domain of graphs) and the reason that robust self-ordering is relevant.²⁶

The query complexity of a tester for Π is a function (of the parameters d, n and ϵ) that represents the number of queries made by the tester on the worst-case n -vertex graph of maximum degree d , when given the proximity parameter ϵ . Fixing d , we typically ignore its effect on the complexity (equiv., treat d as a hidden constant). Also, when stating that the query complexity is $\Omega(q(n))$, we mean that this bound holds for all sufficiently small $\epsilon > 0$; that is, there exists a constant $\epsilon_0 > 0$ such that distinguishing between n -vertex graphs in Π and n -vertex graphs that are ϵ_0 -far from Π requires $\Omega(q(n))$ queries.

5.2 Our first result and the general methodology

With the foregoing preliminaries in place, we state the first result of this section, which is proved using Theorem 4.2.

Theorem 5.2 (linear query complexity lower bound for testing an efficiently recognizable graph property in the bounded-degree graph model): *For any sufficiently large constant d , there exists an efficiently recognizable graph property Π such that testing Π in the bounded-degree graph model (with degree bound d) has query complexity $\Omega(n)$. Furthermore, each n -vertex graph in Π consists of connected components of size $o(\log n)$.*

The main part of the theorem was known before: As observed in [17, Sec. 1], *there exists graph properties that are recognizable in polynomial-time and yet are extremely hard to test in the bounded-degree graph model.* This follows from the fact that the local reduction from testing 3LIN (mod 2) to testing 3-Colorability used by Bogdanov, Obata, and Trevisan [3] is invertible in polynomial-time (which is a common feature of reductions used in the context of NP-completeness proofs).²⁷ Indeed, their reduction actually demonstrates that the set of (3-colorable) graphs that are obtained by applying this reduction to satisfiable 3LIN (mod 2) instances is hard to test (i.e., requires linear query complexity in the bounded-degree graph model).²⁸ We mention that the resulting property contains only connected graphs, which means that Theorem 5.2 has some added value: The fact

²⁵Of course, a tolerant tester is also required to reject with probability at least $2/3$ any graph that is ϵ -far from Π .

²⁶As noted in Section 1.1.1, this is a special case of the general phenomenon pivoted at the difference between ordered and unordered structures, which arises in many contexts (in complexity and logic).

²⁷Of course, 3LIN (i.e., the satisfiability of linear equations (with three variables each) over $\text{GF}(2)$) is easily solvable in polynomial-time. Nevertheless, Bogdanov *et al.* [3] use a reduction of 3LIN to 3-Colorability (via 3SAT) that originates in the theory of NP-completeness in order to reduce between the testing problems.

²⁸Like almost all reductions of this type, the analysis of the reduction actually refers to the promise problem induced by the image of the reduction (i.e., the image of both the yes- and no-instances).

that it applies to graphs with tiny connected components is interesting, since testing properties of such graphs may seem easy (or at least not extremely hard) at first thought.

Proof: Our starting point is a property Φ of (binary) strings (equiv., Boolean functions) that is recognizable in polynomial-time but has a linear query complexity lower bound (see, e.g., [20, Sec. 7]). This refers to a model in which one makes queries to bits of the tested string, and the distance between strings is the (relative) Hamming distance. Such lower bounds were transported to the *dense graph model* in [19, 10.2.3] (see also [20]), but – to the best of own knowledge – no such transportation were performed before in the context of the bounded-degree graph model. Using robustly self-ordered graphs of bounded degree, we present such a transportation.

Construction 5.2.1 (from properties of strings to properties of bounded-degree graphs): *Suppose that $\{G_n = ([n], E_n)\}_{n \in \mathbb{N}}$ is a family of robustly self-ordered graphs of maximum degree $d - 2$.*

- For every $n \in \mathbb{N}$ and $s \in \{0, 1\}^n$, we define the graph $G'_s = ([3n], E'_s)$ such that

$$E'_s = E_n \cup \{\{i, n+i\}, \{i, 2n+i\} : i \in [n]\} \cup \{\{n+i, 2n+i\} : i \in [n] \wedge s_i = 1\} \quad (8)$$

That is, G'_s consists of a copy of G_n augmented by $2n$ vertices such that vertex $i \in [n]$ forms a triangle with $n+i$ and $2n+i$ if $s_i = 1$, and forms a wedge with $n+i$ and $2n+i$ otherwise.

- For a set of strings Φ , we define $\Pi = \bigcup_{n \in \mathbb{N}} \Pi_n$ as the set of all graphs that are isomorphic to some graph G'_s such that $s \in \Phi$; that is,

$$\Pi_n = \{\pi(G'_s) : s \in (\Phi \cap \{0, 1\}^n) \wedge \pi \in \text{Sym}_{3n}\} \quad (9)$$

where Sym_{3n} denote the set of all permutations over $[3n]$.

We may assume, without loss of generality, that G_n has no isolated vertices. Hence, given a graph of the form $\pi(G'_s)$, the vertices of G_n are easily identifiable as having degree at least three (since vertices outside G_n have degree at most two). The foregoing construction yields a local reduction of Φ to Π , where locality means that each query to G'_s can be answered by making a constant number of queries to s , and the (standard) validity of the reduction is based on the fact that G_n is asymmetric.²⁹

In order to be useful towards proving lower bounds on the query complexity of testing Π , we need to show that the foregoing reduction is “distance preserving” (i.e., strings that are far from Φ are transformed into graphs that are far from Π). The hypothesis that G_n is robustly self-ordered is pivotal to showing that if the string s is far from Φ , then the graph G'_s is far from Π .

Claim 5.2.2 (preserving distances): *If $s \in \{0, 1\}^n$ is ϵ -far from Φ , then the $3n$ -vertex graph G'_s (as defined in Construction 5.2.1) is $\Omega(\epsilon)$ -far from Π .*

Proof: We prove the contrapositive. Suppose that G'_s is δ -close to Π . Then, for some $r \in \Phi$ and a permutation $\pi : [3n] \rightarrow [3n]$, it holds that G'_s is δ -close to $\pi(G'_r)$. (The possible use of a non-trivial permutation arises from the fact that Π is closed under isomorphism.) If $\pi(i) = i$ for every $i \in [n]$,

²⁹Standard validity means that $s \in \Phi$ if and only if $G'_s \in \Pi$. Evidently, $s \in \Phi$ is mapped to $G'_s \in \Pi$; the asymmetry of G_n is used to show that $s \notin \Phi$ is mapped to $G'_s \notin \Pi$, since G'_s can not be isomorphic to any graph G'_w such that $w \neq s$. This, by itself, does not mean that if s is far from Φ then G'_s is far from Π .

then s must be $(3d\delta/2)$ -close to r , where d is the degree bound (of the model), since $s_i = 1$ (resp., $r_i = 1$) if and only if i forms a triangle with $n + i$ and $2n + i$ in G'_s (resp., in $\pi(G'_r) = G'_r$).³⁰ Unfortunately, the foregoing condition (i.e., $\pi(i) = i$ for every $i \in [n]$) need not hold in general.

In general, the hypothesis that $\pi(G'_r)$ is δ -close to G'_s implies that π maps at most $3\delta dn/2$ vertices of $[n]$ to $\{n + 1, \dots, 3n\}$. This is the case since each vertex of $[n]$ has degree at least three in G'_r , whereas the other vertices have degree at most two in G'_s (or in any other graph G'_s). Hence, if $t = |\{i \in [n] : \pi(i) \in \{n + 1, \dots, 3n\} \}|$, then $\pi(G'_r)$ and G'_s differ on at least t edges, whereas the hypothesis is that the difference is at most $\delta \cdot 3dn/2$. It follows that $t \leq \delta \cdot 3dn/2$.

Turning to the vertices $i \in [n]$ that π maps to $[n] \setminus \{i\}$, we upper-bound their number by $O(\delta d^2 n)$, since the difference between $\pi(G'_r)$ and G'_s is at most $\delta \cdot 3dn/2$, whereas the hypothesis that G_n is c -robustly self-ordered implies that the difference between $\pi(G'_r)$ and G'_s (or any other graph G'_w) is at least

$$\Delta = c \cdot |\{i \in [n] : \pi(i) \neq i\}| - d \cdot |\{i \in [n] : \pi(i) \notin [n]\}|.$$

(Compare Case 6 in the proof of Theorem 4.2.)³¹

Letting $I = \{i \in [n] : \pi(i) = i\}$, observe that $D \stackrel{\text{def}}{=} |\{i \in I : r_i \neq s_i\}| \leq 3\delta dn/2$, since $r_i \neq s_i$ implies that, for every $i \in I$, the subgraph induced by $\{i, n + i, 2n + i\}$ is different in $\pi(G'_r)$ and G'_s (i.e., it is a triangle in one graph and contains two edges in the other), whereas by the hypothesis $\pi(G'_r)$ and G'_s differ on at most $\delta \cdot 3dn/2$ edges. Recalling that $|I| = n - O(\delta d^2 n)$, it follows that $|\{i \in [n] : r_i \neq s_i\}| \leq (n - |I|) + D = O(\delta d^2 n)$. Recalling that d is a constant, we infer that s is $O(\delta)$ -close to $r \in \Phi$, and the claims follows. ■

Conclusion. Starting with Theorem 4.2 (i.e., an efficient construction of robustly self-ordered graphs of bounded degree), using Construction 5.2.1, and applying Claim 5.2.2, the theorem follows. Specifically, we need to verify the following facts.

- The set Π is polynomial-time recognizable.

Given an $3n$ -vertex graph G' , an adequate algorithm first tries to identify and order the vertices of the corresponding graph G_n , which means that it finds $s \in \{0, 1\}^n$ such that G' is isomorphic to G'_s (or determines that no such s exists). Having found s , the algorithm accepts if and only if $s \in \Phi_n$, where Φ is polynomial-time recognizable by our starting hypothesis.

The first step relies on the hypothesis that G_n can be constructed in polynomial-time, and proceeds as follows.

1. Identifies a set of n vertices, denoted I , in G' such that each vertex in I has degree greater than 2, rejecting if the number of such vertices is different from n .
2. Finds the unique isomorphism between G_n and subgraph of G' induced by I , rejecting if no such isomorphism is found.

³⁰Hence, G'_s is δ -close to G'_r implies that $|\{i \in [n] : s_i \neq r_i\}| \leq \delta \cdot 3dn/2$, which means that s is $\frac{3\delta dn/2}{n}$ -close to r .

³¹Hence, $\Delta \leq \delta \cdot 3dn/2$ implies that

$$\begin{aligned} |\{i \in [n] : \pi(i) \neq i\}| &= \frac{\Delta + d \cdot |\{i \in [n] : \pi(i) \notin [n]\}|}{c} \\ &\leq \frac{3\delta dn/2 + d \cdot 3\delta dn/2}{c} \end{aligned}$$

which is $O(\delta d^2 n)$.

(Here we rely on the fact that isomorphism between graphs of bounded-degree can be found in polynomial-time [30]).

3. The foregoing isomorphism determines the ordering of the vertices in I , which in turn determines s (or indicates that G' is not isomorphic to any G'_s).

- Testing Π requires linear query complexity.

This is shown by reducing testing Φ to testing Π , while recalling that testing Φ requires linear query complexity. Given (proximity parameter ϵ and) oracle access to a string $s \in \{0, 1\}^n$, we invoke the tester for Π (with proximity parameter $\Omega(\epsilon)$) while emulating oracle access to G'_s in a straightforward manner (i.e., each query to G'_s is answered by making at most one query to s). Recall that $s \in \Phi$ implies $G'_s \in \Pi$, whereas by Claim 5.2.2 if s is ϵ -far from Φ then G'_s is $\Omega(\epsilon)$ -far from Π .

This completes the proof, since the n -vertex graphs of Theorem 4.2 have connected components of size $o(\log n)$. ■

Digest: Reducing testing properties of strings to testing graph properties. We wish to highlight the fact that the proof of Theorem 5.2 is based on a general reduction of testing any property Φ of strings to testing a corresponding (bounded-degree) graph property Π . This reduction is described in Construction 5.2.1 and its validity is proved in Claim 5.2.2. Recall that, for any n , the graph property Π consists of $3n$ -vertex graphs (of bounded-degree) that encode the different n -bit long strings in Φ . This reduction is local and preserves distances:

Locality: Each string $s \in \{0, 1\}^n$ is encoded by a graph G'_s such that each query to G'_s can be answered by making at most one query to s .

Preserving distances: If $s \in \Phi$ then $G'_s \in \Pi$, whereas if s is ϵ -far from Φ then G'_s is $\Omega(\epsilon)$ -far from Π .

Recall that G'_s consists of a fixed robustly self-ordered n -vertex graph G_n augmented by (n two-vertex) gadgets that encode s . Let us spell out the effect of this reduction.

Corollary 5.3 (implicit in the proof of Theorem 5.2): *For Φ and Π as in Construction 5.2.1, let Q_Φ and Q_Π denote the query complexities of testing Φ and Π , respectively. Then, $Q_\Phi(n, \epsilon) \leq Q_\Pi(3n, \Omega(\epsilon))$. Likewise, letting Q'_Φ (resp., Q'_Π) denote the query complexity of tolerantly testing Φ (resp., Π), it holds that $Q'_\Phi(n, \eta, \epsilon) \leq Q'_\Pi(3n, \eta/3, \Omega(\epsilon))$.*

The tolerant testing part requires an additional justification. Specifically, we observe that strings s that are η -close to Φ yield graphs G'_s that are $\eta/3$ -close to Π . This is the case because, if the n -bit long strings s and r differ on k bits, then the $3n$ -vertex graphs G'_s and G'_r differ on k vertex pairs.

In preparation to proving the separation between the complexities of testing and tolerant testing, we show a reduction in the opposite direction. This reduction holds provided that the robustly self-ordered graphs used in the definition of Π are locally reversed self-ordered (see Definition 4.8).

Proposition 5.4 (reducing testing Π to testing Φ): *Suppose that the graphs used in Construction 5.2.1 are locally self-ordered and locally reversed self-ordered, and let Φ, Π and Q_Φ, Q_Π be as in*

Corollary 5.3. Then, $Q_{\Pi}(3n, \epsilon) \leq \text{poly}(\log n) \cdot (Q_{\Phi}(n, 2\epsilon) + O(1/\epsilon))$. Furthermore, one-sided error probability is preserved.³²

Recall that the hypothesis can be met by using Corollary 4.10.

Proof: Given oracle access to a graph $G' = ([3n], E')$, we first test that G' is isomorphic to G'_s , for some $s \in \{0, 1\}^n$, and then invoke the tester for Φ while providing it with oracle access to s . Specifically, when the latter tester queries the bit i , we use the local reversed self-order algorithm in order to locate the i^{th} vertex of G_n in G' , and then determine the bit s_i accordingly. Details follow.

Let V denote the set of vertices of the graph $G' = ([3n], E')$ that have degree greater than 2 and neighbor two vertices that have degree at most 2 and neighbor each other if they have degree 2. Evidently, the vertices of V are easy to identify by querying G' for their neighbors and their neighbors' neighbors. Furthermore, $|V| \leq n$, since each vertex in V has two neighbors that are not connected to any other vertex in V , and equality holds in case $G' \in \Pi$. We try to find a (“pivot”) vertex $p \in V$ by picking an arbitrary vertex in G' and checking it and its neighbors. If none of these is in V , then we reject. Otherwise, we continue; we shall be using p as an auxiliary input in all (future) invocations of the local reversed self-ordering algorithm, denoted A .

Intuitively, fixing the pivot p , allow to locate, for each $i \in [n]$, the i^{th} vertex of G_n in G' ; that is, this vertex is $A^{G'}(p, i)$. Indeed, this holds when G' is in Π , but all bets are off otherwise. In particular, when $G' \notin \Pi$, the answer $A^{G'}(p, i)$ may not be a vertex in V , or may equal $A^{G'}(p, j)$ for some $j \neq i$. We avoid these cases by checking that $A^{G'}(p, i) \in V$ and that the vertex in G_n that corresponds to $A^{G'}(p, i)$ is indeed indexed i (by using the local self-ordering algorithm). In particular, we define $A'(i) \stackrel{\text{def}}{=} A^{G'}(p, i)$ if both conditions hold, and let $A'(i)$ be undefined otherwise. Hence, $i \mapsto A'(i)$ is always a bijection from a subset of $[n]$ to a subset of V . This strategy is detailed next.

Using the foregoing algorithm A and the pivot $p \in V$, we define $A'(i) = A^{G'}(p, i)$ if $A^{G'}(p, i) \in V$ and invoking the local self-ordering algorithm on input $A^{G'}(p, i)$ yields i . Otherwise $A'(i)$ is undefined. Hence, evaluating A' amounts to evaluating A as well as evaluating the local self-ordering algorithm. Letting $I' \subseteq [n]$ denote the set of “indices” (i.e., vertices of G_n) on which A' is defined, we note that A' is a bijection from I' to $V' \stackrel{\text{def}}{=} \{A'(i) : i \in I'\} \subseteq V$, and that $I' = [n]$ (and $V' = V$) if $G' \in \Pi$. Hence, our first test is testing whether $I' = [n]$, which is done by selecting at random $O(1/\epsilon)$ elements of $[n]$, and rejecting if A' is undefined on any of them. Otherwise, we proceed, while assuming that $|I'| \geq (1 - 0.1\epsilon) \cdot n$.

Next, we test whether the subgraph of G_n induced by I' is isomorphic to the subgraph of G' induced by V' , where the isomorphism is provided by A' (which maps I' to V'). In other words, we actually test whether the A' -image of the subgraph of G_n induced by I' equals the subgraph of G' induced by V' . This can be done by sampling $O(1/\epsilon)$ vertices of G_n and comparing their neighbors to the neighbors of the corresponding vertices in G' , which are found by A' . Specifically, for every sampled vertex $i \in [n]$, we determine its set of neighbors S_i in G_n , obtain both $A'(i)$ and $A'(S_i) \stackrel{\text{def}}{=} \{A'(j) : j \in S_i\}$, which are supposedly the corresponding vertices in G' , and check whether $A'(S_i)$ is the set of neighbors of $A'(i)$ in G' . We reject if A' is undefined on any of these vertices (i.e., on sampled vertices or their neighbors in G_n). Needless to say, we also reject if any of the foregoing neighborhood checks fails.

³²A tester is said to have one-sided error probability if it always accepts objects that have the property.

Assuming that we did not reject so far, we may assume that G' is $\epsilon/2$ -close to being isomorphic to some G'_s , where the isomorphism is consistent with the inverse of A' . At this point, we invoke the tester for Φ , denoted T , in order to test whether $s \in \Phi$. This is done by providing T with oracle access to s as follows. When T makes a query $i \in [n]$, we determine $A'(i)$, and use our query access to G' in order to determine the two neighbors of $A'(i)$ that have degree at most 2 (and are either connected or have degree 1). If this fails, we reject. Otherwise, we answer 1 if and only if these two neighbors are connected in G' .

To summarize, we employ three tests to G' : An *initial test* of the size I' (which also includes finding a pivot $p \in V$), an *equality test* between the A' -image of the subgraph of G' induced by I' and the subgraph of G_n induced by V' , and an emulation of the testing of Φ . (In all tests, if we encounter an index in $[n] \setminus I'$, we suspend the execution and reject.) For simplicity and without loss of generality, we may assume that T is correct with high (constant) probability.

Note that if $G' \in \Pi$, then it holds that $G' = \pi(G'_s)$ for some $s \in \Phi$ and some permutation $\pi \in \text{Sym}_{3n}$. In this case, it holds that $|I'| = n$ and we always find a pivot $p \in V$. Furthermore, A' equals the restriction of π to $[n]$, the “equality (between induced subgraphs) test” always succeeds, and the emulation of oracle access to s is perfect. Hence, we accept with high probability (or always, if T has one-sided error probability).

On the other hand, suppose that G' is ϵ -far from Π . If either $|I'| < (1 - 0.1\epsilon) \cdot n$ or the subgraph of G' induced by V' is 0.1ϵ -far from the A' -image of the subgraph of G_n induced by I' , then we reject with high probability due to one of the first two tests. Otherwise, letting π be an arbitrary bijection of $[3n]$ to $[3n]$ that extends A' , it follows that for some $s \in \{0, 1\}^n$ the graph G' is 0.2ϵ -close to $\pi(G'_s)$, since we may obtain $\pi(G'_s)$ from G' by modifying the neighborhood of $0.1n$ vertices in I' as well as of the vertices in $[n] \setminus I'$. Furthermore, for every $i \in [n]$ on which A' is defined, it holds that $s_i = 1$ if and only if the two neighbors of $A'(i)$ that have degree at most 2 are connected. By the hypothesis regarding G' , the string s must be $\frac{0.8\epsilon \cdot 3dn/2}{n}$ -far from Φ , and $A'(i) = \pi(i)$ whenever A' is defined on $i \in [n]$. It follows that either the emulation of T was abruptly terminated (leading to rejection) or the answers provided to T are according to s . Hence, we reject with high probability. ■

Digest: Tightly reducing testing properties of strings to testing graph properties. In continuation to (the main part of) Corollary 5.3, we highlight the fact that Construction 5.2.1 not only reduces testing the string property Φ to testing the graph property Π , but rather does so in a rather tight manner. Specifically, for Φ, Π and Q_Φ, Q_Π as in Corollary 5.3, it holds that $Q_\Phi(n, \epsilon)$ and $Q_\Pi(\Theta(n), \Theta(\epsilon))$ agree up to a $\text{poly}(\log n)$ factor. In other words, *for any property of strings Φ , there exists a property of bounded-degree graphs Π such that the (query and time) complexity of testing Φ is reflected in the (query and time) complexity of testing Π* , where our notion of reflection allows for a polylogarithmic slackness. Recall that the transformation of strings in Φ to graphs in Π is (strongly/locally) efficient.

5.3 Separating tolerant testing from testing

Using Corollary 5.3 and Proposition 5.4, we transport the separation between tolerant testing and testing, from the domain of testing strings, where it has been established by [15], to the domain of testing graph properties in the bounded-degree graph model.

Theorem 5.5 (in the bounded-degree graph model, tolerant testing is harder than testing): *For any sufficiently large constant d and any constant $c \in (0, 1)$, there exists a graph property Π such that testing Π in the bounded-degree graph model (with degree bound d) has query complexity $O(\text{poly}(\log n)/\epsilon)$, but tolerantly testing Π has query complexity $\Omega(n^{\Omega(1-c)})$, provided that the tolerance parameter is not smaller than n^{-c} . Furthermore, Π is efficiently recognizable.*

Proof: A small variant on the proof of [15, Thm. 1.3] yields an efficiently recognizable set of strings Φ that is testable in $O(1/\epsilon)$ queries but tolerantly testing it requires $\Omega(n^{\Omega(1-c)})$ queries.³³ Using Construction 5.2.1 with graphs that are locally self-ordered and locally reversed self-ordered (as provided by Corollary 4.10), we obtain the desired graph property Π . By Corollary 5.3 tolerantly testing Π requires $\Omega(n^{\Omega(1)})$ queries, whereas by Proposition 5.4 (non-tolerant) testing Π has query complexity $\text{poly}(\log n) \cdot O(1/\epsilon)$. The claim follows. ■

6 Random Regular Graphs are Robustly Self-Ordered

While Theorem 4.1 only asserts the existence of robustly self-ordered d -regular graphs, we next show that almost all d -regular graphs are robustly self-ordered. This extends work in probabilistic graph theory, which proves a similar result for the weaker notion of self-ordered (a.k.a asymmetric) graphs [4, 5].

Theorem 6.1 (random d -regular graphs are robustly self-ordered): *For any sufficiently large constant d , a random $2d$ -regular n -vertex graph is robustly self-ordered with probability $1 - o(1)$.*

Recall that, with very high probability, these graphs are expanders. We mention that the proof of Theorem 4.1 actually established that n -vertex graphs drawn from a weird distribution (which has min-entropy $\Omega(n)$) are robustly self-ordered with probability $1 - o(1)$. However, this is established by using the edge-coloring variant, and requires employing the transformation presented in Section 2.1. In contrast, the following proof works directly with the original (uncolored) variant, and is completely self-contained.

Proof: The proof is quite similar to the proof Claim 4.1.1, but it faces complications that were avoided in the prior proof by using edge-colors and implicitly directed edges. Specifically, for candidate permutations $\pi_1, \dots, \pi_d : [n] \rightarrow [n]$ (to be used in the construction) and all (non-trivial) permutations $\mu : [n] \rightarrow [n]$, the proof of Claim 4.1.1 considered events of the form $(\forall j \in [d]) \pi_j(i) = \mu(\pi_j(\mu^{-1}(i)))$, whereas here we shall consider events of the form $\{\pi_j^b(i) : j \in [d] \& b \in \{\pm 1\}\} = \{\mu(\pi_k^c(\mu^{-1}(i))) : k \in [d] \& c \in \{\pm 1\}\}$. That is, rather than considering a sequence of equations that refers to single π_j 's, we consider a sequence of equations that refer to the set of all π_j 's, which were expressed above in terms of equations between multi-sets. These multi-set equations will be reduced to equations among sequences by considering all possible ordering of these multi-sets. This amounts to taking a union bound over all possible ordering and results in a more complicated analysis (due to the π_j^{-1} 's) and much more cumbersome notation.

³³Basically, the construction of [15] consists of repeating some m -bit long string $\text{poly}(m)$ times and augmenting it with a PCP of Proximity (PCPP) [2, 11] of membership in some polynomial-time recognizable set that is hard to test. Essentially, the PCPP helps the tester, but it may be totally useless (when corrupted) in the tolerant testing setting. While [15] lets the PCPP occupy an $o(1/\log \log n)$ fraction of the final n -bit string, we let it occupy just a n^{-c} fraction (and use $m = n^{\Omega(1-c)}$). This requires using a different PCPP than the one used in [15]; e.g., using a strong PCPP with linear detection probability [10, Def. 2.2] will do, and such a PCPP is available [10, Thm. 3.3].

To facilitate the proof, we use the standard methodology (cf. [13, Apdx. 2]) of first proving the result in the *random permutation model*, then transporting it to the *configuration model* (by using a general result of [25]), and finally conditioning on the event that the generated graph is simple (which occurs with positive constant probability). Indeed, both models generate multi-graphs that are not necessarily simple graphs (i.e., these multi-graphs may have self-loops and parallel edges). We also use the fact that the simple graphs that are generated by the configuration model (for degree d') are uniformly distributed among all d' -regular graphs.

Recall that in the *random permutation model* a $2d$ -regular n -vertex multi-graph is generated by selecting uniformly and independently d permutations $\pi_1, \dots, \pi_d : [n] \rightarrow [n]$. The multi-graph, denoted $G_{(\pi_1, \dots, \pi_d)}$, consists of the edge multi-set $\bigcup_{j \in [d]} \{\{i, \pi_j(i)\} : i \in [n]\}$, where the $2j^{\text{th}}$ (resp., $(2j - 1)^{\text{st}}$) neighbor of vertex i is $\pi_j(i)$ (resp., $\pi_j^{-1}(i)$). Note that this multi-graph may have self-loops (due to $\pi_j(i) = i$), which contributed two units to the degree of a vertex, as well as parallel edges (due to $\pi_j(i) = \pi_k(i)$ for $j \neq k$ and $\pi_j(i) = \pi_k^{-1}(i)$ for any j, k). We denote the j^{th} neighbor of vertex i by $g_j(i)$; that is, $g_j(i) = \pi_{j/2}(i)$ if j is even, and $g_j(i) = \pi_{(j+1)/2}^{-1}(i)$ otherwise.

Consider an arbitrary permutation $\mu : [n] \rightarrow [n]$, and let $T = \{i \in [n] : \mu(i) \neq i\}$ be its set of non-fixed-point. We shall show that, with probability $1 - \exp(-\Omega(d \cdot |T| \cdot \log n))$ over the choice of $\bar{\pi} = (\pi_1, \dots, \pi_d)$, the size of the symmetric difference between $G_{\bar{\pi}}$ and $\mu(G_{\bar{\pi}})$ is $\Omega(|T|)$. Note that this difference is (half) the sum over $i \in [n]$ of the size of the symmetric difference between the multi-set of neighbors of vertex i in $G_{\bar{\pi}}$ and the multi-set of neighbors of vertex i in $\mu(G_{\bar{\pi}})$. We refer to the latter difference by the phrase *the contribution of vertex i to the difference between $G_{\bar{\pi}}$ and $\mu(G_{\bar{\pi}})$* .

As a warm-up, we first show that each element of T contributes a non-zero number of units to the difference (between $G_{\bar{\pi}}$ and $\mu(G_{\bar{\pi}})$) with probability $1 - O(\text{poly}(d)/n)^{d/6}$ over the choice of $\bar{\pi}$. Recalling that $i \in T$ contributes to the difference (between $G_{\bar{\pi}}$ and $\mu(G_{\bar{\pi}})$) if the multi-sets of its neighbors in $G_{\bar{\pi}}$ and $\mu(G_{\bar{\pi}})$ differ, it follows that $i \in T$ contributes to the difference if *for every permutation $\sigma : [2d] \rightarrow [2d]$ there exists $j \in [2d]$ such that $g_j(i) \neq \mu(g_{\sigma(j)}(\mu^{-1}(i)))$* . Hence, the complementary event holds (i.e., i does not contribute to the difference) if and only if there exists permutation $\sigma : [2d] \rightarrow [2d]$ such that for every $j \in [2d]$ it holds that $g_j(i) = \mu(g_{\sigma(j)}(\mu^{-1}(i)))$. Thus, the probability that i does not contribute to the difference is given by

$$\begin{aligned} & \Pr_{\bar{\pi}} [\exists \sigma \in \text{Sym}_{2d} (\forall j \in [2d]) \ g_j(i) = \mu(g_{\sigma(j)}(\mu^{-1}(i)))] \\ & \leq (2d)! \cdot \max_{\sigma \in \text{Sym}_{2d}} \{ \Pr_{\bar{\pi}} [(\forall j \in [2d]) \ g_j(i) = \mu(g_{\sigma(j)}(\mu^{-1}(i)))] \}. \end{aligned} \quad (10)$$

Fixing σ that maximizes the probability, and denoting it σ_i , consider any $J_i \subseteq [d]$ such that for the j 's in J_i the multi-sets $\{j, \lceil \sigma_i(2j)/2 \rceil\}$'s are disjoint (i.e., $\{j, \lceil \sigma_i(2j)/2 \rceil\} \cap \{k, \lceil \sigma_i(2k)/2 \rceil\} = \emptyset$ for any $j \neq k \in J_i$). Note that we may select J_i such that $|J_i| \geq d/6$, since taking j to J_i rules out taking (to J_i) at most five other values (i.e., letting $v \stackrel{\text{def}}{=} \lceil \sigma_i(2j)/2 \rceil$, the value k is ruled out if either $k \in \{j, v\}$ or $\sigma_i(2k) \in \{2j - 1, 2j, 2v - 1, 2v\}$). Using this feature of J_i , we prove –

Claim 6.1.1 (warm-up):³⁴ Eq. (10) is upper-bounded by $(2d)^{2d} \cdot (3/n)^{|J_i|}$.

³⁴One may obtain a better bound of $O(d/n)^{2d}$ by analyzing Eq. (10) directly, by considering all the $2d$ events and accounting for their small dependency. On the other hand, we can obtain higher robustness parameter by considering smaller sets J_i 's (say of size $d/12$), which suffice for counting vertices that contribute (say) $d/2$ units to the difference between $G_{\bar{\pi}}$ and $\mu(G_{\bar{\pi}})$.

Proof: Recalling that σ_i maximizes Eq. (10), we upper-bound Eq. (10) by

$$\begin{aligned} & (2d)! \cdot \Pr_{\bar{\pi}} [(\forall j \in J_i) \ g_{2j}(i) = \mu(g_{\sigma_i(2j)}(\mu^{-1}(i)))] \\ &= (2d)! \cdot \prod_{j \in J_i} \Pr_{\pi_j, \pi_{\lceil \sigma_i(2j)/2 \rceil}} [g_{2j}(i) = \mu(g_{\sigma_i(2j)}(\mu^{-1}(i)))] \end{aligned} \quad (11)$$

where the equality uses the disjointness of the multi-sets $\{j, \lceil \sigma_i(2j)/2 \rceil\}$ for the j 's in J_i . Next, letting $\sigma'_i(2j) \stackrel{\text{def}}{=} \lceil \sigma_i(2j)/2 \rceil$, and $\sigma''_i(2j) \stackrel{\text{def}}{=} (-1)^{\sigma_i(2j) \bmod 2}$ (so that $\sigma_i(2j) = 2\sigma'_i(2j) - 0.5(1 - \sigma''_i(2j))$), we upper-bound Eq. (11) by

$$(2d)! \cdot \prod_{j \in J_i} \Pr_{\pi_j, \pi_{\sigma'_i(2j)}} [\pi_j(i) = \mu(\pi_{\sigma'_i(2j)}^{\sigma''_i(2j)}(\mu^{-1}(i)))] < (2d)^{2d} \cdot (3/n)^{|J_i|}, \quad (12)$$

where $\Pr_{\pi_j, \pi_j}[\cdot]$ stands for $\Pr_{\pi_j}[\cdot]$ and π^1 stands for π , while the inequality is justified by considering the following three cases (w.r.t each $j \in J_i$).

1. If $k \stackrel{\text{def}}{=} \sigma'_i(2j) \neq j$ (equiv., $\sigma_i(2j) \notin \{2j-1, 2j\}$), then, letting $b = \sigma''_i(2j)$, the corresponding factor in the l.h.s of Eq. (12) is

$$\Pr_{\pi_j, \pi_k} [\pi_j(i) = \mu(\pi_k^b(\mu^{-1}(i)))]$$

which equals $1/n$ by fixing π_k , letting $v = \mu(\pi_k^b(\mu^{-1}(i)))$, and using $\Pr_{\pi_j}[\pi_j(i) = v] = 1/n$.

2. If $\sigma_i(2j) = 2j$, then the corresponding factor in the l.h.s of Eq. (12) is

$$\Pr_{\pi_j} [\pi_j(i) = \mu(\pi_j(\mu^{-1}(i)))]$$

which is at most $1/(n-1)$ since $\mu(i) \neq i$; specifically, fixing the value of $\pi_j(\mu^{-1}(i))$, and denoting this value by v , leaves $\pi_j(i)$ uniformly distributed in $[n] \setminus \{v\}$, which means that $\Pr_{\pi_j}[\pi_j(i) = \mu(v) \mid v = \pi_j(\mu^{-1}(i))] \leq 1/(n-1)$ (where equality holds if $\mu(v) \neq v$).

3. If $\sigma_i(2j) = 2j-1$, then the corresponding factor in the l.h.s of Eq. (12) is

$$\Pr_{\pi_j} [\pi_j(i) = \mu(\pi_j^{-1}(\mu^{-1}(i)))]$$

which is shown to be less than $3/n$. In this case, we consider two sub-cases depending on whether or not $\pi_j(i) = \mu^{-1}(i)$, while noting that the first case occurs with probability $1/n$ whereas $\Pr_{\pi_j}[\pi_j(i) = \mu(\pi_j^{-1}(\mu^{-1}(i))) \mid \pi_j(i) \neq \mu^{-1}(i)] \leq 1/(n-1)$.

Hence, each of the factors in the l.h.s of Eq. (12) is upper-bounded by $3/n$, and the claim follows. \blacksquare

The general case. The same argument generalizes to any set $I \subseteq T$ such that $I \cap \mu(I) = \emptyset$. In such a case we get

$$\begin{aligned} & \Pr_{\bar{\pi}} [(\forall i \in I) (\exists \sigma_i \in \text{Sym}_{2d}) (\forall j \in [2d]) \ g_j(i) = \mu(g_{\sigma_i(j)}(\mu^{-1}(i)))] \\ & \leq (2d)!^{|I|} \cdot \max_{\sigma_1, \dots, \sigma_n} \left\{ \Pr_{\bar{\pi}} [(\forall i \in I) (\forall j \in [2d]) \ g_j(i) = \mu(g_{\sigma_i(j)}(\mu^{-1}(i)))] \right\} \end{aligned} \quad (13)$$

Claim 6.1.2 (actual analysis): Eq. (13) is upper-bounded by

$$(2d)^{2d \cdot |I|} \cdot (3/(n - 2(|I| - 1)))^{|I| \cdot d/6}. \quad (14)$$

Proof: Fixing $\sigma_1, \dots, \sigma_n$ that maximize Eq. (13), we proceed as in the proof of Claim 6.1.1. For every $i \in I = \{i_1, \dots, i_m\}$, we fix a set J_i of size at least $d/6$ such that the multi-sets $\{j, \lceil \sigma_i(2j)/2 \rceil\}$'s are disjoint, and refer to the events $E_{j,k}(\pi_1, \dots, \pi_{2d})$ that depend only on the value of π_j and $\pi_{\sigma'_{i_k}(2j)}$ on the points i_1, \dots, i_{k-1} and $\mu^{-1}(i_1), \dots, \mu^{-1}(i_{k-1})$, respectively, where $\sigma'_i(2j) \stackrel{\text{def}}{=} \lceil \sigma_i(2j)/2 \rceil$, and $\sigma''_i(2j) \stackrel{\text{def}}{=} (-1)^{\sigma_i(2j) \bmod 2}$ (as in the proof of Claim 6.1.1). Specifically, $E_{j,k}(\pi_1, \dots, \pi_{2d})$ is the event

$$(\forall k' \in [k-1]) \quad g_{2j}(i_{k'}) = \mu(g_{\sigma_{i_{k'}}(2j)}(\mu^{-1}(i_{k'})))$$

which can be written as

$$(\forall k' \in [k-1]) \quad \pi_j(i_{k'}) = \mu(\pi_{\sigma'_{i_{k'}}(2j)}(\mu^{-1}(i_{k'}))).$$

Using the disjointness of the $\{j, \lceil \sigma_i(2j)/2 \rceil\}$'s in J_i , we upper-bound Eq. (13) by

$$\begin{aligned} & (2d)!^m \cdot \prod_{k \in [m]} \prod_{j \in J_{i_k}} \Pr_{\pi_1, \dots, \pi_{2d}} \left[g_{2j}(i_k) = \mu(g_{\sigma_{i_k}(2j)}(\mu^{-1}(i_k))) \mid E_{j,k}(\pi_1, \dots, \pi_{2d}) \right] \\ &= (2d)!^m \cdot \prod_{k \in [m]} \prod_{j \in J_{i_k}} \Pr_{\pi_1, \dots, \pi_{2d}} \left[\pi_j(i_k) = \mu(\pi_{\sigma'_{i_k}(2j)}(\mu^{-1}(i_k))) \mid E_{j,k}(\pi_1, \dots, \pi_{2d}) \right]. \end{aligned} \quad (15)$$

Now, when analyzing the foregoing conditional probability in Eq. (15), we consider two cases for each $k \in [m]$ and $j \in J_{i_k}$. If $j \neq \sigma'_{i_k}(2j)$, then we fix the value of each of these two permutations (i.e., π_j and $\pi_{\sigma'_{i_k}(2j)}$) on the corresponding $k-1$ points that occur in the condition $E_{j,k}$, and the value of these permutations on the k^{th} points (i.e., i_k and $\mu^{-1}(i_k)$) is restricted accordingly (i.e., to the remaining $n - (k-1)$ values). Otherwise (i.e., $j = \sigma'_{i_k}(2j)$), we fix the value of π_j on these $2(k-1)$ points. Hence, the argument in the warm-up analysis applies with n replaces by either $n - (k-1)$ or $n - 2(k-1)$. It follows that Eq. (15) is upper-bounded by

$$(2d)!^m \cdot \prod_{k \in [m]} (3/(n - 2(m-1)))^{|J_{i_k}|}.$$

Using $|J_{i_k}| \geq d/6$ for every $k \in [m]$, the claim follows. ■

Recall that Eq. (14) refers to a fixed set $I \subseteq T$ such that $I \cap \mu(I) = \emptyset$, and that it constitutes an upper bound on the probability (over the choice of $\bar{\pi}$) that, for each $i \in I$ there exists a permutation $\sigma_i : [2d] \rightarrow [2d]$ such that $g_j(i) = \mu(g_{\sigma_i(j)}(\mu^{-1}(i)))$ holds for all $j \in [2d]$. This upper bound (i.e., $(2d)^{2d \cdot |I|} \cdot (3/(n - 2(|I| - 1)))^{|I| \cdot d/6}$) simplifies to $(2d)^{2d \cdot |I|} \cdot (5/n)^{|I| \cdot d/6}$, provided that $|I| \leq \lceil n/6 \rceil$.

Recalling that $t \stackrel{\text{def}}{=} |T| \in [n]$, we shall upper-bound the probability (over the choice of $\bar{\pi}$) that T contains a $\lceil t/2 \rceil$ -subset T' such that for each $i \in T'$ there exists a permutation $\sigma_i : [2d] \rightarrow [2d]$ such that $g_j(i) = \mu(g_{\sigma_i(j)}(\mu^{-1}(i)))$ holds for all $j \in [2d]$. We do so by taking a union bound over all $\lceil t/6 \rceil$ -subsets I such that $I \cap \mu(I) = \emptyset$ and for each $i \in I$ there exists a permutation $\sigma_i : [2d] \rightarrow [2d]$

such that $g_j(i) = \mu(g_{\sigma_i(j)}(\mu^{-1}(i)))$ holds for all $j \in [2d]$. (Note that such a $\lceil t/6 \rceil$ -subset I exists in each $\lceil t/2 \rceil$ -subset T' , and that $\lceil t/6 \rceil < n/3$.) Using the aforementioned simplified form of Eq. (14), we conclude that, with probability at most

$$\binom{t}{\lceil t/6 \rceil} \cdot (2d)^{2d \cdot \lceil t/6 \rceil} \cdot (5/n)^{\lceil t/6 \rceil \cdot d/6} < 2^t \cdot (5 \cdot (2d)^{12}/n)^{\lceil t/6 \rceil \cdot d/6} = \exp(-\Omega(dt \log n))$$

over the choice of $\bar{\pi}$, the set T contains no $\lceil t/6 \rceil$ -subset I as above. This means that, with probability at most $\exp(-\Omega(dt \log n))$, less than $t/2$ of the indices $i \in T$ contribute a non-zero number of units to the difference (between $G_{\bar{\pi}}$ and $\mu(G_{\bar{\pi}})$).

Letting $c' = 1/2$ and considering all (non-trivial) permutations $\mu : [n] \rightarrow [n]$, we conclude that the probability, over the choice of $\bar{\pi}$, that $G_{\bar{\pi}}$ is not c' -robustly self-ordered is at most

$$\begin{aligned} \sum_{t \in [n]} \binom{n}{t} \cdot \exp(-\Omega(dt \log n)) &= \sum_{t \in [n]} \exp(-\Omega((d - O(1)) \cdot t \log n)) \\ &= \exp(-\Omega((d - O(1)) \cdot \log n)), \end{aligned}$$

and the claim follows for the permutation model (and for any sufficiently large d).

As stated upfront, using the general result of [25, Thm. 1.3], we infer that a uniformly distributed $2d$ -regular n -vertex multi-graph fails to be c' -robustly self-ordered with probability $o(1)$. Lastly, recalling that such a $2d$ -regular multi-graph is actually a simple graph with probability $\exp(-((2d)^2 - 1)/4)$, the theorem follows. ■

Digest. The proof of Theorem 6.1 is quite similar to the proof Claim 4.1.1, but it faces two complications that were avoided in the prior proof (by using edge-colors and implicitly directed edges). Most importantly, the current proof has to handle equality between multi-sets instead of equality between sequences. This is done by considering all possible ordering of these multi-sets, which amounts to taking a union bound over all possible ordering and results in more complicated analysis and notation. (Specifically, see the introduction of σ_i 's and J_i 's and the three cases analyzed in the warm-up.) In addition, since edges are defined by permutations over the vertex-set rather than by perfect matching, we have to consider both the forward and backward direction of each permutation, which results in further complicating the analysis and the notation. (Specifically, see the introduction of σ'_i 's and σ''_i 's and the three cases analyzed in the warm-up.)

An alternative proof of Theorem 4.2. We mention that combining an extension of Theorem 6.1 with some of the ideas underlying the proof of Theorem 4.2 yields an alternative proof of Theorem 4.2 (i.e., an alternative construction of robustly self-ordered bounded-degree graphs).

Remark 6.2 (an alternative construction of d -regular robustly self-ordered graphs): *On input 1^n , we set $\ell = \frac{O(\log n)}{\log \log n}$, and proceeds in three steps.*

1. *Extending the proof of Theorem 6.1, we show that for all sufficiently large constant d , for any set \mathcal{G} of $t = t(\ell) < n = \ell^{\Omega(\ell)}$ ($2d$ -regular) ℓ -vertex graphs, with probability $1 - o(1)$, a random $2d$ -regular ℓ -vertex graph is both robustly self-ordered and far from being isomorphic to any graph in \mathcal{G} . Note that, with probability $1 - o(1)$, such a graph is also expanding.*

Here two ℓ -vertex graphs are said to be far apart if they disagree on $\Omega(\ell)$ vertex-pairs.

The proof of Theorem 6.1 is extended by considering, for a random graph, the event that it is either not robustly self-ordered or is not far from an isomorphic copy of one of the t (fixed) graphs. The later event (i.e., being close to isomorphic to one of these graphs) occurs with probability $o(t/n)$.

2. Relying on Step 1, we find a sequence of n/ℓ robustly self-ordered $2d$ -regular ℓ -vertex graphs that are expanding and pairwise far from being isomorphic to one another.

This is done by iteratively finding robustly self-ordered $2d$ -regular ℓ -vertex expanding graphs that are far from being isomorphic to all prior ones, where scanning all possible graphs and checking the condition can be done in time $n \cdot \ell^{d\ell/2} \cdot (\ell!) = \text{poly}(n)$.

3. Using the sequence of n/ℓ graphs found in Step 2, we consider the n -vertex graph that consists of these ℓ -vertex graphs as its connected components, and use parts of the proof of Theorem 4.2 to show that this graph is robustly self-ordered. Specifically, we only need to consider cases that are analogous to Cases 2, 6 and 7. The treatment of the analogous cases is slightly simpler than in the proof of Theorem 4.2, since the graphs are somewhat simpler.

Note that the resulting graphs are not locally constructable.

Part II

The Case of Dense Graphs

Recall that when considering graphs of unbounded degree, we ask whether we can obtain unbounded robustness parameters. In particular, we are interested in n -vertex graphs that are $\Omega(n)$ -robustly self-ordered, which means that they must have $\Omega(n^2)$ edges.

In Section 7 we prove the existence of $\Omega(n)$ -robustly self-ordered n -vertex graphs, and show that they imply $\Omega(1)$ -robustly self-ordered bounded-degree $O(n^2)$ -vertex graphs. In Section 8, we reduce the construction of the former (dense) n -vertex graphs to the construction of non-malleable two-source extractors (with very mild parameters). We actually show two reductions: The first reduction (presented in Section 8.1) requires the extractors to have an additional natural feature, called quasi-orthogonality, and yields a construction of such n -vertex graphs that runs in $\text{poly}(n)$ -time. The second reduction (presented in Section 8.2) does not make this additional requirement, and yields an algorithm that computes the adjacency predicate of such n -vertex graphs in $\text{poly}(\log n)$ -time.

In Section 9 we demonstrate the applicability of $\Omega(n)$ -robustly self-ordered n -vertex graphs to property testing; specifically, to proving lower bounds (on the query complexity) for the dense graph testing model. Lastly, in Section 10, we consider the construction of $\Omega(d(n))$ -robustly self-ordered n -vertex graphs of maximum degree $d(n)$, for every $d : \mathbb{N} \rightarrow \mathbb{N}$ such that $d(n) \in [\Omega(1), n]$.

7 Existence and Transformation to Bounded-Degree Graphs

It seems easier to prove that random n -vertex graphs are $\Omega(n)$ -robustly self-ordered (see Proposition 7.1) than to prove that random bounded-degree graphs are $\Omega(1)$ -robustly self-ordered (or even just prove that such bounded-degree graphs exist). In contrast, it seems harder to construct

$\Omega(n)$ -robustly self-ordered n -vertex graphs than to construct $\Omega(1)$ -robustly self-ordered bounded-degree graphs. In particular, we show that $\Omega(n)$ -robustly self-ordered n -vertex graphs can be easily transformed into $O(n^2)$ -vertex bounded-degree graphs that are $\Omega(1)$ -robustly self-ordered (see Proposition 7.2). We stress that the construction of robustly self-ordered bounded-degree graphs that is obtained by combining the foregoing transformation with Theorem 1.4 is entirely different from the constructions presented in the first part of the paper.

Random graphs are robustly self-ordered. We first show that, with very high probability, a random n -vertex graph $G_n = ([n], E_n)$, where E_n is a uniformly distributed subset of $\binom{[n]}{2}$, is $\Omega(n)$ -robustly self-ordered.

Proposition 7.1 (robustness analysis of a random graph): *A random n -vertex graph $G_n = ([n], E_n)$ is $\Omega(n)$ -robustly self-ordered with probability $1 - \exp(-\Omega(n))$.*

As stated above, the following proof is significantly easier than the proof provided for the bounded-degree analogue (i.e., Theorem 6.1).

Proof: For each (non-trivial) permutation $\mu : [n] \rightarrow [n]$, letting $T \stackrel{\text{def}}{=} \{i \in [n] : \mu(i) \neq i\}$ denote its (non-empty) set of non-fixed-points, we show that, with probability $1 - \exp(-\Omega(n \cdot |T|))$, the size of the symmetric difference between a random n -vertex graph $G_n = ([n], E_n)$ and $\mu(G_n)$ is $\Omega(n \cdot |T|)$.

For every $u, v \in [n]$ such that $u < v$, let $\chi_{u,v} = \chi_{u,v}^\mu(G_n)$ represent the event that *the pair $(\mu(u), \mu(v))$ contributes to the symmetric difference between G_n and $\mu(G_n)$* ; that is, $\chi_{u,v} = 1$ if exactly one of the edges $\{\mu(u), \mu(v)\}$ and $\{u, v\}$ is in G_n , since $\{u, v\}$ is an edge of G_n if and only if $\{\mu(u), \mu(v)\}$ is an edge of $\mu(G_n)$. Note that $\Pr[\chi_{u,v}(G_n) = 1] = 1/2$ if $\{u, v\} \neq \{\mu(u), \mu(v)\}$, because fixing the adjacency relation of the pair $(\mu(u), \mu(v))$ leaves the adjacency relation of pair (u, v) totally random. We shall prove that

$$\Pr_{G_n} \left[\sum_{u < v \in [n]} \chi_{u,v}^\mu(G_n) < \frac{n \cdot |T|}{20} \right] = \exp(-\Omega(n \cdot |T|)). \quad (16)$$

We prove Eq. (16) by identifying a set S of $\Omega(n \cdot |T|)$ vertex pairs such that the random variables in $\{\chi_{u,v}^\mu(G_n) : \{u, v\} \in S\}$ are totally independent and uniformly distributed in $\{0, 1\}$. Specifically if S and $S' = \{\{\mu(u), \mu(v)\} : \{u, v\} \in S\}$ are disjoint, then fixing the adjacencies of the pairs in S' leaves the adjacency of pairs in S totally random.

Fixing a $\lceil |T|/3 \rceil$ -subset $I \subseteq T$ such that $I \cap \mu(I) = \emptyset$, let $J = [n] \setminus (I \cup \mu^{-1}(I))$. Note that $\mu(I \cup J) \subseteq [n] \setminus I$ and that $|J| = n - 2 \cdot \lceil |T|/3 \rceil \geq (n/3) - 2$. Observe that, for every $(u, v) \in J \times I$, it holds that $u \neq v$ and $\Pr[\chi_{u,v} = 1] = 1/2$, where the equality is due to $\{u, v\} \neq \{\mu(u), \mu(v)\}$, which holds since $(u, v) \in J \times I$ but $\mu(u), \mu(v) \in [n] \setminus I$. Furthermore, the events the correspond to the pairs in $J \times I$ are independent, because the sets $\{\{u, v\} : (u, v) \in J \times I\}$ and $\{\{\mu(u), \mu(v)\} : (u, v) \in J \times I\}$ are disjoint; that is, $(u, v) \in J \times I$ implies $(\mu(u), \mu(v)) \in ([n] \setminus I) \times ([n] \setminus I)$. Hence (using $n \leq 3(|J| + 2)$ and $|T| \leq 3|I|$ (as well as $3(|J| + 2) \cdot 3|I| < 9.9 \cdot |J| \cdot |I|$)), the l.h.s. of Eq. (16) is upper-bounded by

$$\begin{aligned} \Pr_{G_n} \left[\sum_{(u,v) \in J \times I} \chi_{u,v}^\mu(G_n) < \frac{3(|J| + 2) \cdot 3|I|}{20} \right] &\leq \Pr_{G_n} \left[\sum_{(u,v) \in J \times I} \chi_{u,v}^\mu(G_n) < \frac{0.99 \cdot |J| \cdot |I|}{2} \right] \\ &= \exp(-\Omega(|J| \cdot |I|)) \end{aligned}$$

which is $\exp(-\Omega(n \cdot |T|))$. Having established Eq. (16), the claim follows by a union bound (over all non-trivial permutations $\mu : [n] \rightarrow [n]$); specifically, denoting the set of non-trivial permutations by P_n , we upper-bound the probability that G_n is not $\frac{n}{20}$ -robust by

$$\begin{aligned} & \sum_{\mu \in P_n} \Pr_{G_n} [\mu \text{ violates the condition in Eq. (16)}] \\ & \leq \sum_{t \in [n]} \binom{n}{t} \cdot (t!) \cdot \exp(-\Omega(n \cdot t)) \\ & < n \cdot \max_{t \in [n]} \{n^t \cdot \exp(-\Omega(n \cdot t))\} \\ & = \exp(-\Omega(n)) \end{aligned}$$

where t represents the size of the set of non-fixed-points (w.r.t μ). ■

Obtaining bounded-degree robustly self-ordered graphs. We next show how to transform $\Omega(n)$ -robustly self-ordered n -vertex graphs to $O(n^2)$ -vertex bounded-degree graphs that are $\Omega(1)$ -robustly self-ordered. Essentially, we show that the standard “degree reduction via expanders” technique works (when using a different color for the expanders’ edges, and then using gadgets to replace colored edges). Specifically, we replace each vertex in $G_n = ([n], E_n)$ by an $(n - 1)$ -vertex expander graph and connect each of these vertices to at most one vertex in a different expander, while coloring the edges of the expanders with 1, and coloring the other edges by 2. Actually, the vertex v is replaced by the vertex-set $C_v = \{\langle v, u \rangle : u \in [n] \setminus \{v\}\}$ and in addition to the edges of the expander, colored 1, we connect each vertex $\langle v, u \rangle \in C_v$ to the vertex $\langle u, v \rangle \in C_u$ and color this edge 2 if $\{u, v\} \in E_n$ and 0 otherwise.³⁵ This yields an $n \cdot (n - 1)$ -vertex $O(1)$ -regular graph, denoted G'_n , coupled with an edge-coloring, denoted χ' , which uses three colors. We stress that each vertex in G'_n is incident to a single even-colored edge. Using the hypothesis that G_n is $\Omega(n)$ -robustly self-ordered, we prove that (G'_n, χ') is $\Omega(1)$ -robustly self-ordered (in the colored sense).

Proposition 7.2 (robustness analysis of the degree reduction): *If G_n is $\Omega(n)$ -robustly self-ordered, then (G'_n, χ') is $\Omega(1)$ -robustly self-ordered (in the colored sense of Definition 2.1).*

Using Theorem 2.4 (after adding self-loops), we obtain a $O(1)$ -regular $O(n^2)$ -vertex graph that is $\Omega(1)$ -robustly self-ordered (in the standard sense).

Proof: Denoting the vertex-set of G'_n by $V = \bigcup_{v \in [n]} C_v$, we consider an arbitrary (non-trivial) permutation $\mu' : V \rightarrow V$, and the corresponding set of non-fixed-points T' . Intuitively, if μ' maps vertices of C_v to several C_w ’s, then we get a proportional contribution to the difference between G'_n and $\mu'(G'_n)$ by the (1-colored) edges of the expander. Otherwise, μ' induces a permutation μ over the vertices of G_n , and we get a corresponding contribution via the (2-colored) edges of G_n . Lastly, non-identity mappings inside the individual C_v ’s are charged using the (even-colored) edges that connect different C_v ’s, while relying on the fact each vertex in G'_n is incident to a single even-colored edge. Details follow.

For a permutation $\mu' : V \rightarrow V$ as above, let $\mu : [n] \rightarrow [n]$ be a permutation that maximizes the (average over $v \in [n]$ of the) number of vertices in C_v that are mapped by μ' to vertices in $C_{\mu(v)}$;

³⁵This is equivalent to first converting G_n into a n -vertex clique while coloring an edge 2 if and only if it is in E_n .

that is, for every permutation $\nu : [n] \rightarrow [n]$, it holds that

$$|\{\langle v, u \rangle \in V : \mu'(\langle v, u \rangle) \in C_{\mu(v)}\}| \geq |\{\langle v, u \rangle \in V : \mu'(\langle v, u \rangle) \in C_{\nu(v)}\}|. \quad (17)$$

We consider the following three cases.

Case 1: $\sum_{v \in [n]} |B_v| = \Omega(|T'|)$, where $B_v \stackrel{\text{def}}{=} \{\langle v, u \rangle \in C_v : \mu'(\langle v, u \rangle) \notin C_{\mu(v)}\}$.

(This refers to the case that many vertices are mapped by μ' to an expander that is different from the one designated by μ , which represents the best possible mapping of whole expanders.)

Letting $C_{v,w} \stackrel{\text{def}}{=} \{\langle v, u \rangle : \mu'(\langle v, u \rangle) \in C_w\}$, we first observe that for every v it holds that $\max_{w \neq \mu(v)} |C_{v,w}| \leq \frac{2}{3} \cdot (n-1)$, because otherwise we reach a contradiction to the maximality of μ by defining $\nu(v) = w$ and $\nu(\mu^{-1}(w)) = \mu(v)$, where w is the element obtaining the maximum, and $\nu(x) = \mu(x)$ otherwise.

Next, observe that there exists $W_v \subseteq [n] \setminus \{\mu(v)\}$ such that $B'_v = \bigcup_{w \in W_v} C_{v,w}$ satisfies both $|B'_v| \leq \frac{2}{3} \cdot (n-1)$ and $|B'_v| \geq |B_v|/3$. Now, consider the sets B'_v and $C_v \setminus B'_v$: On the one hand, in $\mu'(G'_n)$ there are $\Omega(|B'_v|)$ 1-colored edges connecting $\mu'(B'_v)$ and $\mu'(C_v \setminus B'_v)$, due to the subgraph of $\mu'(G'_n)$ induced by $\mu'(C_v)$ which equals subgraph of G'_n induced by C_v (which, in turn, is an expander). On the other hand, in G'_n there are no 1-colored edges between $\mu'(B'_v)$ and $\mu'(C_v \setminus B'_v)$, since $\mu'(B'_v) \subseteq \bigcup_{w \in W_v} C_w$ and $\mu'(C_v \setminus B'_v) \subseteq \bigcup_{w \in [n] \setminus W_v} C_w$.

We conclude that, in this case, the difference between G'_n and $\mu'(G_n)$ is $\sum_v \Omega(|B'_v|) = \sum_v \Omega(|B_v|) = \Omega(|T'|)$.

Case 2: $\sum_{v \in [n]; \mu(v) \neq v} |C'_v| = \Omega(|T'|)$, where $C'_v \stackrel{\text{def}}{=} \{\langle v, u \rangle \in C_v : \mu'(\langle v, u \rangle) \in C_{\mu(v)}\}$.

(This refers to the case that many vertices are mapped by μ' to an expander that is designated by μ , but this expander is not the one in which they reside (i.e., μ has many non-fixed-points).)

Letting $\gamma > 0$ be a constant such that G_n is $\gamma \cdot n$ -robustly self-ordered, we may assume that $\sum_{v \in [n]; \mu(v) \neq v} |C'_v| \geq (1 - 0.5 \cdot \gamma) \cdot \sum_{v \in [n]; \mu(v) \neq v} |C_v|$, since otherwise we are done by Case 1.

By the γn -robust self-ordering of G_n , the difference between G_n and $\mu(G_n)$ is at least $\Delta \stackrel{\text{def}}{=} \gamma n \cdot |\{v \in [n] : \mu(v) \neq v\}|$. Assuming, for a moment, that $\mu'(C_v) = C_v$ for every v such that $\mu(v) \neq v$, the difference between G'_n and $\mu'(G'_n)$ is Δ , where the difference is due to edges colored 2 (i.e., the edges inherited from G_n). This amount is proportional to the number of vertices in the current case, since

$$\Delta = \frac{\gamma n}{n-1} \cdot \sum_{v: \mu(v) \neq v} |C_v| > \gamma \cdot \sum_{v: \mu(v) \neq v} |C_v|.$$

In general, $\mu'(C_v) = C_v$ may not hold for some v , and in this case we may lose the contribution of the 2-colored edges incident at vertices in $\bigcup_{v \in [n]; \mu(v) \neq v} (C_v \setminus C'_v)$. Recalling that (by our hypothesis) the size of this set is at most $0.5 \cdot \gamma \cdot \sum_{v: \mu(v) \neq v} |C_v|$, we are left with a contribution of at least $0.5\gamma \cdot \sum_{v: \mu(v) \neq v} |C'_v|$.

We conclude that, in this case, the difference between G'_n and $\mu'(G_n)$ is $\Omega(\sum_{v: \mu(v) \neq v} |C'_v|) = \Omega(|T'|)$.

Case 3: $\sum_{v \in [n]} |C_v''| = \Omega(|T'|)$, where $C_v'' \stackrel{\text{def}}{=} \{\langle v, u \rangle \in C_v : \mu'(\langle v, u \rangle) \in C_v \setminus \{\langle v, u \rangle\}\}$.

(This refers to the case that many vertices are mapped by μ' to a different vertex in the same expander in which they reside.)³⁶

(This case would have been easy to handle if the expanders used on the C_v 's were robustly self-ordered. Needless to say, we want to avoid such an assumption. Instead, we rely on the fact that in G_n' different vertices in C_v are connected to different C_u 's.)

We may assume that $\sum_{v \in [n]} |C_v''| \geq 2 \cdot \sum_{v \in [n]} |\{\langle v, u \rangle \in C_v : \mu'(\langle v, u \rangle) \notin C_v\}|$, since otherwise we are done by either Case 1 or Case 2 (see Footnote 36). Now, consider a generic $\langle v, u \rangle \in C_v''$, and let $w \neq u$ be such that $\mu'(\langle v, u \rangle) = \langle v, w \rangle$. Then, in $\mu'(G_n')$ an edge colored either 0 or 2 connects $\langle v, w \rangle = \mu'(\langle v, u \rangle)$ to $\mu'(\langle u, v \rangle)$, since $\langle v, u \rangle$ and $\langle u, v \rangle$ are so connected in G_n' , whereas in G_n' an even-colored edge connects $\langle v, w \rangle$ to $\langle w, v \rangle \in C_w$. Recall, however, that $\langle v, w \rangle$ is incident to a single even-colored edge. We consider two sub-cases.

- If $\mu'(\langle u, v \rangle) \in C_u$, then $\langle v, w \rangle$ contributes to the difference between $\mu'(G_n')$ and G_n' , because in $\mu'(G_n')$ vertex $\langle v, w \rangle$ is connected (by its unique even-colored edge) to a vertex in C_u whereas in G_n' vertex $\langle v, w \rangle$ is connected (by its unique even-colored edge) to a vertex in C_w .

(Recall that w is uniquely determined by $\langle v, u \rangle \in C_v''$, since $\mu'(\langle v, u \rangle) = \langle v, w \rangle$, and so this contribution can be charged to $\langle v, u \rangle$.)

- If $\mu'(\langle u, v \rangle) \notin C_u$, then $\langle u, v \rangle$ resides in the set $\bigcup_{x \in [n]} \{\langle x, y \rangle \in C_x : \mu'(\langle x, y \rangle) \notin C_x\}$, which (by the hypothesis) has size at most $0.5 \cdot \sum_{v \in [n]} |C_v''|$

Hence, at least half of $\bigcup_{v \in [n]} C_v''$ appears in the first sub-case, which implies that, in this case, the difference between G_n' and $\mu'(G_n)$ is at least $\frac{1}{2} \cdot \sum_{v \in [n]} |C_v''| = \Omega(|T'|)$.

Hence, the difference between G_n' and $\mu'(G_n)$ is $\Omega(|T'|)$. ■

8 Relation to Non-Malleable Two-Source Extractors

For $n = 2^\ell$, we reduce the construction of $\Omega(n)$ -robustly self-ordered (dense) n -vertex graphs to the construction of non-malleable two-source extractors for $(\ell, \ell - O(1))$ -sources. Recall that a random variable X is called an (ℓ, k) -source if X is distributed over $[2^\ell]$ and has min-entropy at least k (i.e., $\Pr[X=i] \leq 2^{-k}$ for every $i \in [2^\ell]$).³⁷ A function $\mathbf{E} : [2^\ell] \times [2^\ell] \rightarrow \{0, 1\}^m$ is called a (standard) two-source (k, ϵ) -extractor if, for every two independent (ℓ, k) -sources X and Y , it holds that $\mathbf{E}(X, Y)$ is ϵ -close to the uniform distribution over $\{0, 1\}^m$, denoted U_m . Our notion of a non-malleable two-source extractor, presented next, is one of (the weakest of) the notions considered in [8, 7].

Definition 8.1 (non-malleable two-source extractors [7, Def. 1.3]): *A function $\text{nmE} : [2^\ell] \times [2^\ell] \rightarrow \{0, 1\}^m$ is called a non-malleable two-source (k, ϵ) -extractor if, for every two independent (ℓ, k) -sources X and Y , and for every two functions $f, g : [2^\ell] \rightarrow [2^\ell]$ such that at least one of them has*

³⁶Note that if $\langle v, u \rangle \in C_v$ is not mapped by μ' to C_v , then either $\mu'(\langle v, u \rangle) \notin C_{\mu(v)}$ holds (i.e., Case 1) or $\mu'(\langle v, u \rangle) \in C_{\mu(v)}$ such that $\mu(v) \neq v$ (i.e., Case 2). Hence, if $\langle u, v \rangle \in T'$ is not counted in Cases 1 and 2, then it must be counted in Case 3.

³⁷Indeed, for the sake of simplicity (of our arguments), we do not require that $\ell \in \mathbb{N}$, but rather only that $2^\ell \in \mathbb{N}$; consequently, we consider distributions over $[2^\ell]$ rather than over $\{0, 1\}^\ell$.

no fixed-points, it holds that $(\mathbf{nmE}(X, Y), \mathbf{nmE}(f(X), g(Y)))$ is ϵ -close to $(U_m, \mathbf{nmE}(f(X), g(Y)))$; that is,

$$\frac{1}{2} \cdot \sum_{\alpha, \beta} |\Pr[(\mathbf{nmE}(X, Y), \mathbf{nmE}(f(X), g(Y))) = (\alpha, \beta)] - 2^{-m} \cdot \Pr[\mathbf{nmE}(f(X), g(Y)) = \beta]| \leq \epsilon. \quad (18)$$

The parameter ϵ is called the error of the extractor.

We shall be interested in the special case in which f and g are permutations. In this case, the foregoing condition (i.e., Eq. (18)) can be replaced by requiring that $(\mathbf{nmE}(X, Y), \mathbf{nmE}(f(X), g(Y)))$ is 2ϵ -close to the uniform distribution over $\{0, 1\}^{m+m}$.³⁸ Furthermore, we shall focus on non-malleable two-source (k, ϵ) -extractors that output a single bit (i.e., $m = 1$), and in this case non-triviality mandates $\epsilon < 0.5$. In general, we view ϵ as a constant, but view ℓ and k as varying (or generic) parameters, and focus on the case of $k = \ell - O(1)$.

Recall that constructions of non-malleable two-source (k, ϵ) -extractors with much better parameters are known [7, Thm. 1]. In particular, these constructions support $k = \ell - \ell^{\Omega(1)}$, negligible error (i.e., $\epsilon = \exp(-\ell^{\Omega(1)})$), and $m = \ell^{\Omega(1)}$. We stress that, as is the norm in the context of randomness extraction, the extracting function is computable in polynomial-time (i.e., in $\text{poly}(\ell)$ -time).

We shall show that any non-malleable two-source $(\ell - O(1), 0.49)$ -extractor (for sources over $[2^\ell]$) yields $\Omega(2^\ell)$ -robustly self-ordered $O(2^\ell)$ -vertex graphs. Actually, we shall show two such constructions: The first construction runs in $\text{poly}(2^\ell)$ -time, and the second construction provides strong constructability (a.k.a local computability) as claimed in Theorem 1.4. Both constructions use a similar underlying reasoning, which is more transparent in the first construction.

8.1 The first construction

For the first construction, we need the extractor to satisfy the following natural (and quite minimal) requirement, which we call quasi-orthogonality. We say that an extractor $\mathbf{nmE} : [2^\ell] \times [2^\ell] \rightarrow \{0, 1\}$ is quasi-orthogonal (with error ϵ) if the following conditions hold:

1. *The residual function obtained from \mathbf{nmE} by any fixing of one of its two arguments is almost unbiased:* For every $x \in [2^\ell]$ and every $\sigma \in \{0, 1\}$ it holds that $|\{y \in [2^\ell] : \mathbf{nmE}(x, y) = \sigma\}| \leq (0.5 + \epsilon) \cdot 2^\ell$; ditto for every $y \in [2^\ell]$ and the corresponding set $\{x \in [2^\ell] : \mathbf{nmE}(x, y) = \sigma\}$.
2. *The residual functions obtained from \mathbf{nmE} by any two different fixings of one of its two arguments are almost uncorrelated:* For every $\{x, x'\} \in \binom{[2^\ell]}{2}$ it holds that $|\{y \in [2^\ell] : \mathbf{nmE}(x, y) \neq \mathbf{nmE}(x', y)\}| \geq (0.5 - \epsilon) \cdot 2^\ell$; ditto for every $\{y, y'\} \in \binom{[2^\ell]}{2}$ and the corresponding set $\{x \in [2^\ell] : \mathbf{nmE}(x, y) \neq \mathbf{nmE}(x, y')\}$.

As shown in Proposition 8.2, any non-malleable two-source (k, ϵ) -extractor can be transformed (in $\text{poly}(2^\ell)$ -time) to a quasi-orthogonal one at a cost of a small degradation in the parameters (i.e., ϵ increases by an additive term of $O(2^{-(\ell-k)})$ and 2^ℓ decreases by an additive term of $O(2^k)$). Note that $\text{poly}(2^\ell)$ -time is acceptable when one aims at constructing $O(2^\ell)$ -vertex graphs; however, aiming at strong/local constructability (as in Theorem 1.4), we shall avoid such a transformation in the second construction (presented in Section 8.2).

³⁸In this case, $f(X)$ and $g(Y)$ have min-entropy at least k , which implies that $\mathbf{nmE}(f(X), g(Y))$ is ϵ -close to the uniform distribution over $\{0, 1\}^m$.

Proposition 8.2 (transforming non-malleable two-source extractors into ones that are quasi-orthogonal): *For every $k \leq \ell - 3$, there exists a poly(2^ℓ)-time transformation that given a non-malleable two-source (k, ϵ) -extractor $\text{nmE} : [2^\ell] \times [2^\ell] \rightarrow \{0, 1\}$, returns a non-malleable two-source (k, ϵ') -extractor $\text{nmE}' : [n'] \times [n'] \rightarrow \{0, 1\}$ such that $n' \geq 2^\ell - O(2^k)$ and nmE' is quasi-orthogonal with error $\epsilon' = \epsilon + O(2^k/n')$.*

Proof: Essentially, nmE' is obtained from nmE by simply discarding inputs that violate the quasi-orthogonality conditions. Letting $n = 2^\ell$, first note that the number of x 's that violate the first condition is at most 2^{k+1} , because otherwise we obtain a contradiction to the hypothesis that nmE is a two-source (k, ϵ) -extractor (by letting X be uniform on the x 's that satisfy $|\{y \in [n] : \text{nmE}(x, y) = \sigma\}| > (0.5 + \epsilon) \cdot n$ for either $\sigma = 0$ or $\sigma = 1$, and Y be uniform on $[n]$). Next, consider the residual (k, ϵ) -extractor $\text{nmE}_1 : [n_1] \times [n_1] \rightarrow \{0, 1\}$, where $n_1 \geq n - 2^{k+1}$, obtained by omitting the exceptional x 's. Note that nmE_1 satisfies the first quasi-orthogonality condition with respect to the first argument with error ϵ . Doing the same for the second argument yields a residual (k, ϵ) -extractor $\text{nmE}_2 : [n_2] \times [n_2] \rightarrow \{0, 1\}$, where $n_2 \geq n_1 - 2^{k+1}$ and nmE_2 satisfies the first quasi-orthogonality condition (for both arguments) with error $\epsilon + \frac{2^{k+1}}{n_1}$. Likewise, we claim that there are at most 2^k disjoint pairs $\{x, x'\}$'s that violate the second condition (i.e., $|\{y \in [n_2] : \text{nmE}_2(x, y) \neq \text{nmE}_2(x', y)\}| \geq (0.5 - \epsilon) \cdot n_2$), because otherwise we obtain a contradiction to the hypothesis that nmE_2 is a *non-malleable* two-source (k, ϵ) -extractor (by using a function f that maps each such x to its matched x' , and the identity permutation g).³⁹ And, again, we consider a residual extractor obtained by omitting the exceptional pairs. Doing the same for the y 's, we obtained the desired extractor. ■

Recall that non-malleable two-source extractors with much stronger parameters than we need (i.e., min-entropy $\ell - \ell^{\Omega(1)}$, negligible error, and $\ell^{\Omega(1)}$ bits of output), were constructed in [7, Thm. 1], but these extractors are not quasi-orthogonal. Employing Proposition 8.2, we obtain a quasi-orthogonal non-malleable two-source $(\ell - 4, 0.1)$ -extractor that can be used in the construction of the following Theorem 8.3. Essentially, the construction consists of a bipartite graph, with 2^ℓ vertices on each side, such that the edges between the two sides are determined by the extractor. In addition, we add a clique on one of the two sides so that the two sides are (robustly) distinguishable (see Figure 1). We stress that the resulting $2^{\ell+1}$ -vertex graph is $\Omega(2^\ell)$ -robustly self-ordered as long as the non-malleable extractor is quasi-orthogonal and works for very mild parameters; that is, we only require error that is bounded away from $1/2$ with respect to min-entropy $\ell - O(1)$.

Theorem 8.3 (using a quasi-orthogonal non-malleable two-source extractor to obtain a $\Omega(2^\ell)$ -robustly self-ordered $O(2^\ell)$ -vertex graph): *For a constant $\epsilon \in (0, 0.5)$ varying $\ell \geq k$ such that $k \leq \ell - 2 + \log_2(0.5 - \epsilon) = \ell - O(1)$, suppose that $\text{nmE} : [2^\ell] \times [2^\ell] \rightarrow \{0, 1\}$ is a quasi-orthogonal (with error ϵ) non-malleable two-source (k, ϵ) -extractor. Then, the $2^{\ell+1}$ -vertex graph $G = (V_1 \cup V_0, E)$ such that $V_\sigma = \{\langle \sigma, i \rangle : i \in [2^\ell]\}$ and*

$$E = \{\{\langle 1, i \rangle, \langle 0, j \rangle\} : \text{nmE}(i, j) = 1\} \cup \binom{V_1}{2} \quad (19)$$

is $\Omega(|V_1 \cup V_0|)$ -robustly self-ordered. Furthermore, the claim holds even if the non-malleability condition (i.e., Eq. (18)) holds only for permutations f and g that have no fixed-points.

³⁹Formally, denoting the exceptional pairs by (x_i, x'_i) , where $i \in [2^k]$, we define f such that $f(x_i) = x'_i$ for each $i \in [2^k]$, and let X be uniform on $\{x_i : i \in [2^k]\}$ and Y be uniform on $[n_2]$. Then, $\Pr[\text{nmE}_2(X, Y) \neq \text{nmE}_2(f(X), g(Y))] < 0.5 - \epsilon$, in contradiction to the hypothesis regarding nmE_2 .

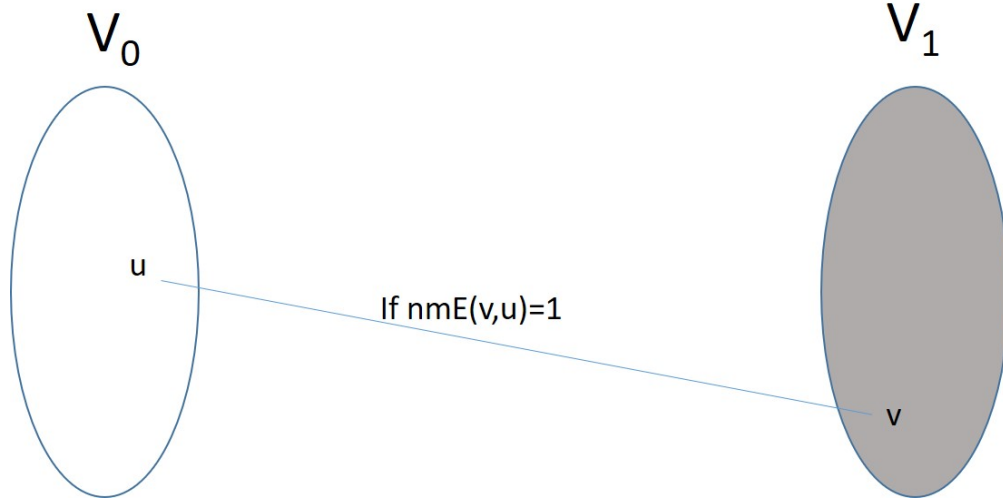


Figure 1: Illustrating the construction of Theorem 8.3.

Indeed, the first set of edges, denoted E' , corresponds to a bipartite graph between V_1 and V_0 that is determined by nmE , and the second set corresponds to a 2^ℓ -vertex clique (see Figure 1). Note that the extraction parameters are extremely weak; that is, the min-entropy may be very high (i.e., $k = \ell - O(1)$), the error may be an arbitrary non-trivial constant (i.e., $\epsilon < 1/2$), and we only extract one bit (i.e., $m = 1$).

Proof: Let $V = V_1 \cup V_0$, and consider an arbitrary (non-trivial) permutation $\mu : V \rightarrow V$. Intuitively, if μ maps a vertex of V_1 to V_0 , then the difference in degrees of vertices in the two sets (caused by the clique edges) contributes at least $(2^\ell - 1) - 2\epsilon \cdot 2^\ell$ units to the symmetric difference between G and $\mu(G)$, where here we use the first quasi-orthogonality condition. On the other hand, if μ maps $\langle 1, i \rangle \in V_1$ to $V_1 \setminus \{\langle 1, i \rangle\}$, then the difference in the neighborhoods caused by the bipartite graph contributes at least $(0.5 - \epsilon) \cdot 2^\ell$ units to the symmetric difference between G and $\mu(G)$. To prove this, we distinguish between the case that μ has relatively few non-fixed-points (in either V_0 or V_1), which is analyzed using the second quasi-orthogonality condition, and the case that μ has relatively many non-fixed-points (in both V_0 and V_1), which is analyzed using the non-malleability condition. Details follow.

Let $T = \{v \in V : \mu(v) \neq v\}$ denote the set of non-fixed-points of μ . Then, we consider two types of vertices: Those that belong to the set $T' = \bigcup_{\sigma \in \{0,1\}} \{v \in V_\sigma : \mu(v) \notin V_\sigma\} \subseteq T$ and those that belong to $T \setminus T'$. The threshold for distinguishing these cases is set to $K = (0.5 - \epsilon) \cdot 2^{\ell-2} = \Omega(|V|)$.

Case 1: $|T'| \geq K$.

(This refers to the case that many vertices are mapped by μ to the opposite side of the bipartite graph (V, E') , where ‘many’ means $\Omega(|V|)$.)

Each vertex in T' contributes $(1 - 2\epsilon) \cdot 2^\ell - 1$ units to the symmetric difference between G and $\mu(G)$, because the degree of each vertex in V_1 is at least $(2^\ell - 1) + (0.5 - \epsilon) \cdot 2^\ell$, whereas the degree of each vertex in V_0 is at most $(0.5 + \epsilon) \cdot 2^\ell$, where we use the first quasi-orthogonality condition, which implies that the number of bipartite edges incident at each vertex is at least $(0.5 - \epsilon) \cdot 2^\ell$ and at most $(0.5 + \epsilon) \cdot 2^\ell$.

Hence, the symmetric difference between G and $\mu(G)$ is at least $((1 - 2\epsilon) \cdot 2^\ell - 1) \cdot |T'| = \Omega(|V|) \cdot |T'|$, since $2^\ell = \Omega(|V|)$. Using the case's hypothesis, we have $|T'| = \Omega(|V|) = \Omega(|T|)$, which means that in this case the difference between G and $\mu(G)$ is $\Omega(|V|) \cdot |T|$.

We stress that the difference between G and $\mu(G)$ is at least $\Omega(|V|) \cdot |T'|$ also if the case hypothesis does not hold.

Case 2: $|T'| < K$.

(This refers to the case that few vertices are mapped by μ to the opposite side of the bipartite graph (V, E') , where ‘few’ means less than $K \leq |V|/20$ (assuming $\epsilon \leq 0.1$).)

For every $\sigma \in \{0, 1\}$, let $V'_\sigma = V_\sigma \cap \mu^{-1}(V_\sigma)$ and $T_\sigma = V'_\sigma \cap T$. Indeed, (T', T_0, T_1) is a three-way partition of T . Note that the size of the symmetric difference between G and $\mu(G)$ is lower-bounded by

$$|\{(v, u) \in V'_1 \times V'_0 : \mathbf{nmE}(\mu(v), \mu(u)) \neq \mathbf{nmE}(v, u)\}|, \quad (20)$$

since, for any $(v, u) \in V'_1 \times V'_0$, it holds that $\mu(v)$ neighbors $\mu(u)$ in G if and only if $\mathbf{nmE}(\mu(v), \mu(u)) = 1$, whereas $\mu(v)$ neighbors $\mu(u)$ in $\mu(G)$ if and only if v neighbors u in G which holds if and only if $\mathbf{nmE}(v, u) = 1$.

We consider two sub-cases according to whether or not $\min(|T_0|, |T_1|)$ is relatively large. The threshold for distinguishing these sub-cases is also set to $K = (0.5 - \epsilon) \cdot 2^{\ell-2}$; note that $K = \Omega(|V|)$ and $K \geq 2^k$.

Case 2.1: $\min(|T_0|, |T_1|) < K$.

In this case we shall use the (second condition of) quasi-orthogonality of \mathbf{nmE} .

Suppose, without loss of generality, that $|T_0| \leq |T_1|$, which implies $|T_0| < K$. Then, the contribution of each vertex $v \in T_1$ to Eq. (20) equals

$$\begin{aligned} & |\{u \in V'_0 : \mathbf{nmE}(\mu(v), \mu(u)) \neq \mathbf{nmE}(v, u)\}| \\ & \geq |\{u \in V'_0 : \mathbf{nmE}(\mu(v), u) \neq \mathbf{nmE}(v, u)\}| - |T_0| \\ & \geq |\{u \in V_0 : \mathbf{nmE}(\mu(v), u) \neq \mathbf{nmE}(v, u)\}| - |T'| - |T_0| \\ & \geq (0.5 - \epsilon) \cdot 2^\ell - 2 \cdot K \\ & = (0.5 - \epsilon) \cdot 2^{\ell-1} \end{aligned}$$

where the first inequality uses $\mu(u) = u$ for $u \in V'_0 \setminus T_0$, the second inequality uses $|V'_0| \geq |V_0| - |T'|$, the third inequality uses $\mu(v) \neq v$ along with the (second condition of) quasi-orthogonality of \mathbf{nmE} (and the hypotheses regarding $|T'|$ and $|T_0|$), and the equality is due to $K = (0.5 - \epsilon) \cdot 2^{\ell-2}$.

Hence, in this case, the total contribution to Eq. (20) is $(0.5 - \epsilon) \cdot 2^{\ell-1} \cdot |T_1|$, which is $\Omega(|V|) \cdot (|T| - |T'|)$, since $|T_1| \geq (|T| - |T'|)/2$.

Case 2.2: $\min(|T_0|, |T_1|) \geq K$.

In this case we shall use the non-malleable feature of \mathbf{nmE} .

Specifically, for each $\sigma \in \{0, 1\}$, let μ_σ denote the restriction of μ to T_σ . Essentially, using $K \geq 2^k$, the non-malleability condition of the (k, ϵ) -extractor \mathbf{nmE} implies

$$|\{(i, j) \in T_0 \times T_1 : \mathbf{nmE}(i, j) \neq \mathbf{nmE}(\mu_0(i), \mu_1(j))\}| \geq (0.5 - \epsilon) \cdot |T_0| \cdot |T_1|.$$

This can be seen by letting X and Y be uniform over T_0 and T_1 , respectively. (Also, for sake of formality, we extend μ_0 and μ_1 (which have no fixed-points) to (arbitrary) derangements f and g , respectively.)⁴⁰ In this case, the non-malleability condition implies that the distribution $(\text{nmE}(X, Y), \text{nmE}(f(X), g(Y)))$ is ϵ -close to $(U_1, \text{nmE}(f(X), g(Y)))$, which implies that $\Pr[\text{nmE}(X, Y) \neq \text{nmE}(\mu_0(X), \mu_1(Y))]$ is at least $\Pr[U_1 \neq \text{nmE}(\mu_0(X), \mu_1(Y))] - \epsilon = 0.5 - \epsilon$.

Hence, in this case, the total contribution to Eq. (20) is $(0.5 - \epsilon) \cdot |T_0| \cdot |T_1| = \Omega(|V|) \cdot (|T| - |T'|)$, where we use $\min(|T_0|, |T_1|) = \Omega(|V|)$ and $|T_0| + |T_1| = |T| - |T'|$.

Hence, in both sub-cases, the difference between G and $\mu(G)$ is $\Omega(|V|) \cdot (|T| - |T'|)$.

Recall that (by the last comment at Case 1) the difference between G and $\mu(G)$ is $\Omega(|V|) \cdot |T'|$. Combining this lower-bound with the conclusion of Case 2, the difference between G and $\mu(G)$ is $\Omega(|V|) \cdot |T|$. ■

Digest: Note that the quasi-orthogonality of nmE was used in Cases 1 and 2.1, whereas the non-malleability of nmE (w.r.t derangements) was used in Case 2.2. In particular, Case 1 only uses the first condition of quasi-orthogonality, and does so in order to infer that the degrees of all vertices in the bipartite graph are approximately equal. In Case 2.1 the second quasi-orthogonality condition is used in order to assert that the neighborhoods of two different vertices in V_σ are significantly different. This is useful only when the number of non-fixed-points in $V_{1-\sigma}$ is relatively small. When the number of non-fixed-points in both V_σ 's is large but few vertices are mapped to the other side (i.e., $|T'| \ll |T|$), we only use Case 2.2, which does not refer to quasi-orthogonality at all. Hence, we have the following –

Remark 8.4 (a special case of Theorem 8.3): *For bipartite graphs $G = (V, E)$ such that $V = V_0 \cup V_1$ and $E \subseteq V_0 \times V_1$, we consider the special case of robust self-ordering that refers only to permutations $\mu : V \rightarrow V$ that are derangements that preserve the bipartition of V (i.e., μ has no fixed-points and $\mu(V_0) = V_0$).⁴¹ In this case, assuming (only) that nmE is a non-malleable two-source (ℓ, ϵ) -extractor (i.e., the case of $k = \ell$), implies that, for any such μ , the size of the symmetric difference between G and $\mu(G)$ is $(0.5 \pm \epsilon) \cdot |V_0| \cdot |V_1|$. Furthermore, the claim holds even if the non-malleability condition holds only for permutations f and g that have no fixed-points (i.e., derangements), and the quasi-orthogonality condition is not necessary. Note that the proof of Theorem 8.3 simplifies, since $T' = \emptyset$ and $T_\sigma = V_\sigma = V'_\sigma$ hold, and the size of the symmetric difference between G and $\mu(G)$ equal the quantity in Eq. (20).*

Interestingly, the special case of Theorem 8.3 asserted in Remark 8.4 can be reversed in the sense that a bipartite graph that is robustly self-ordered in the foregoing restricted sense is actually a non-malleable two-source $(\ell, 0.5 - \Omega(1))$ -extractor (w.r.t derangement). In the following result ϵ is an arbitrary constant (in $(0, 0.5)$), whereas G varies and $o(1)$ vanishes with $|V|$.

Proposition 8.5 (a reversal of the special case of Theorem 8.3 (i.e., of Remark 8.4)): *Let $G = (V_0 \cup V_1, E)$ be a bipartite graph such that $|V_0| = |V_1|$ and $E \subseteq V_0 \times V_1$. Let $V = V_0 \cup V_1$, and*

⁴⁰Note that we may assume, w.l.o.g., that $|T_\sigma \cup \mu(T_\sigma)| \leq |V_\sigma| - 2$.

⁴¹That is, the requirement regarding the symmetric difference between G and $\mu(G)$ is made only for permutations μ that have no fixed-points and satisfy $\mu(V_0) = V_0$.

suppose that for every derangement $\mu : V \rightarrow V$ such that $\mu(V_0) = V_0$ it holds that the size of the symmetric difference between G and $\mu(G)$ is $(0.5 \pm \epsilon) \cdot |V_0| \cdot |V_1|$. Then, $F : V_0 \times V_1 \rightarrow \{0, 1\}$ such that $F(x, y) = 1$ if and only if $\{x, y\} \in E$ is a non-malleable two-source $(\ell, \epsilon + \sqrt{2\epsilon} + o(1))$ -extractor w.r.t derangement.⁴²

Needless to say, the claim holds also if G is augmented by complete graph on the vertex-set V_1 . Note that we lose a $\sqrt{2\epsilon} + o(1)$ term in the reversal.

Proof: Let (f, g) and (X, Y) be as in Definition 8.1, and note that in this case X and Y are independent distributions that are each uniformly distributed on $[2^\ell]$. Define $\mu : V \rightarrow V$ such that $\mu(z) = f(z)$ if $z \in V_0$ and $\mu(z) = g(z)$ otherwise, and note that μ is a derangement that preserves the partition of V . Recall that $(\mu(x), \mu(y))$ contributes to the symmetric difference between G and $\mu(G)$ if and only if $F(\mu(x), \mu(y)) \neq F(x, y)$, since $\mu(x)$ is connected to $\mu(y)$ in $\mu(G)$ if and only if x is connected to y in G . Hence, by the hypothesis, we have

$$\Pr[F(X, Y) \neq F(\mu(X), \mu(Y))] = 0.5 \pm \epsilon. \quad (21)$$

Letting $p_{\sigma, \tau}^\mu \stackrel{\text{def}}{=} \Pr[(F(X, Y), F(\mu(X), \mu(Y))) = (\sigma, \tau)]$, we have $p_{0,1}^\mu + p_{1,0}^\mu = 0.5 \pm \epsilon$, and using the fact that (X, Y) and $(\mu(X), \mu(Y))$ are identically distributed we have $p_{1,0}^\mu = p_{0,1}^\mu$ (since $p_{1,1}^\mu + p_{1,0}^\mu = p_{1,1}^\mu + p_{0,1}^\mu$). Hence, $p_{0,1}^\mu = 0.25 \pm 0.5\epsilon$. Lastly, we show that $p_{1,1}^\mu + p_{1,0}^\mu = 0.5 \pm \sqrt{\epsilon/2} + o(1)$, and conclude that $p_{1,1}^\mu = 0.25 \pm (0.5\epsilon + \sqrt{\epsilon/2} + o(1))$; it follows that F is a non-malleable (two-source) $(\ell, \epsilon + \sqrt{2\epsilon} + o(1))$ -extractor.

To show that $p_{1,1}^\mu + p_{1,0}^\mu = 0.5 \pm \sqrt{\epsilon/2} + o(1)$, we first note that $p \stackrel{\text{def}}{=} p_{1,1}^\mu + p_{1,0}^\mu = \Pr[F(X, Y) = 1]$ is actually oblivious of μ . Hence, by considering a random derangement μ that preserves V_0 (i.e., $\mu(V_0) = V_0$), we observe that, with overwhelmingly high probability (over the choice of μ), it holds that $\{(x, y) \in V_0 \times V_1 : F(x, y) \neq F(\mu(x), \mu(y))\}$ has size $(2p(1-p) \pm o(1)) \cdot |V_0| \cdot |V_1|$. Confronting this with Eq. (21), we infer that $p = 0.5 \pm (\sqrt{\epsilon/2} + o(1))$. ■

Corollary. Combining Theorem 8.3 with the non-malleable two-source extractors of [7, Thm. 1], while using Proposition 8.2, we obtain an efficient construction of $\Omega(n)$ -robustly self-ordered graphs (alas not a strongly explicit (aka locally computable) one).

Theorem 8.6 (constructing $\Omega(n)$ -robustly self-ordered n -vertex graphs): *There exist an algorithm that, on input n , works in $\text{poly}(n)$ -time and outputs an explicit description of an $\Omega(n)$ -robustly self-ordered $O(n)$ -vertex graph. Furthermore, each vertex in this graph has degree at least $0.24 \cdot n$ and at most $0.76 \cdot n$.*

The degree bounds follow by observing that the vertices in the graph described in Theorem 8.3 have degree at least $(0.5 - \epsilon) \cdot n/2$ and at most $(1.5 + \epsilon) \cdot n/2$, whereas [7, Thm. 1] provides for $\epsilon = o(1)$.

⁴²That is, the non-malleability condition (i.e., Eq. (18)) is guaranteed only for permutations f and g that have no fixed-points.

8.2 The second construction

Combining Theorem 8.3 with the non-malleable two-source extractors of [7, Thm. 1], while using Proposition 8.2, we obtained an efficient construction of $\Omega(n)$ -robustly self-ordered n -vertex graphs (see Theorem 8.6). Unfortunately, this construction is not locally computable (as postulated in Theorem 1.4), because the non-malleable two-source extractors of [7, Thm. 1] are not quasi-orthogonal and the transformation of Proposition 8.2 runs in time that is polynomial in the size of the resulting graph.

To avoid the foregoing transformation and prove Theorem 1.4, we employ a variant on the construction presented in Theorem 8.3. Rather than connecting two sets of vertices using a bipartite graph that corresponds to a *quasi-orthogonal* non-malleable two-source extractor, we connect three sets of vertices such that one pair of vertex-sets is connected by a (not necessarily quasi-orthogonal) non-malleable two-source extractor, whereas the other two pairs are connected by bipartite graphs that are merely quasi-orthogonal. In analogy to the definition of a quasi-orthogonal (two-source) extractor, we say that a bipartite graph on the vertex-set $X \cup Y$ is **quasi-orthogonal** (with error ϵ) if the following two conditions hold regarding its adjacency predicate $B : X \times Y \rightarrow \{0, 1\}$:

1. *The degree of each vertex is approximately half the number of the vertices on the other side:* For each $x \in X$ (resp., $y \in Y$), it holds that $|\{y \in Y : B(x, y) = 1\}| = (0.5 \pm \epsilon) \cdot |Y|$ (resp., $|\{x \in X : B(x, y) = 1\}| = (0.5 \pm \epsilon) \cdot |X|$).
2. *For each pair of vertices on one side, approximately half the vertices on the other side neighbor exactly one of the vertices in the pair:* For every $x \neq x' \in X$, it holds that $|\{y \in Y : B(x, y) \neq B(x', y)\}| = (0.5 \pm \epsilon) \cdot |Y|$. Similarly, for $y \neq y' \in Y$.

We note that inner-product (mod 2) extractor, denoted $E_2 : \{0, 1\}^\ell \times \{0, 1\}^\ell \rightarrow \{0, 1\}$, corresponds to a quasi-orthogonal bipartite graph for the case $X = Y = \{0, 1\}^\ell \setminus \{0^\ell\}$. We will however need quasi-orthogonal bipartite graphs with different-sized sides, which can be obtained by a simple variant. Specifically, for the case of $X = \{0, 1\}^\ell \setminus \{0^\ell\}$ and $Y = \{0, 1\}^{\ell+2} \setminus \{0^{\ell+2}\}$, we use the function $B(x, y) = E_2(G(x), y)$, where $G : \{0, 1\}^\ell \rightarrow \{0, 1\}^{\ell+2}$ is a small-bias generator that satisfies $G(x) \neq 0^{\ell+2}$ and $G(x) \neq G(x')$ for every $x \neq 0^\ell$ and $x' \neq x$ (see Proposition 8.8, and note that $G(a, b, c, d) = (a, b, c, d, E_2(a, b), E_2(c, d))$ will do). We stress that the foregoing construction is strongly explicit (i.e., locally computable).

We shall also assume that the (bipartite graph corresponding to the) non-malleable extractor $\text{nmE} : [2^\ell - 1] \times [2^\ell - 1] \rightarrow \{0, 1\}$ has *linear degrees* in the sense that for every x it holds that $|\{y \in [2^\ell - 1] : \text{nmE}(x, y) = 1\}| \geq \epsilon' \cdot 2^\ell$ for some constant $\epsilon' > 0$. This can be enforced by starting with an arbitrary non-malleable two-source (k, ϵ') -extractor (e.g., the one of [7, Thm. 1]) and resetting pairs in $m = \epsilon' \cdot 2^\ell$ fixed perfect matchings to 1 (i.e., for each (x, y) in one of these matching, we reset $\text{nmE}(x, y) \leftarrow 1$).⁴³ This increases the error of the extractor by an additive term of $m/2^k = 2^{\ell-k} \cdot \epsilon'$, which we can afford (e.g., $\epsilon' = 0.01$ and $k = \ell - 4$, yields extraction error $\epsilon < 0.2$). We stress that this transformation preserves polynomial-time computability of the extracting function.

Theorem 8.7 (using a non-malleable two-source extractor with linear degrees to obtain a $\Omega(2^\ell)$ -robustly self-ordered $O(2^\ell)$ -vertex graph): *For any constants $\epsilon, \epsilon' \in (0, 0.5)$ and varying $k \leq \ell - 4$,*

⁴³For example, we may use the matchings $\{(z, z + i) : z \in [2^\ell - 1]\}$ for $i \in [m]$, where addition is mod $2^\ell - 1$. Furthermore, starting from an extractor that is defined over ℓ -bit strings, we may omit one of these strings (and obtain an extractor defined over $[2^\ell - 1]$).




Figure 2: Illustrating the construction of Theorem 8.7.

where $\ell \in \mathbb{N}$, suppose that $\mathbf{nmE} : [2^\ell - 1] \times [2^\ell - 1] \rightarrow \{0, 1\}$ is a non-malleable two-source (k, ϵ) -extractor such that for every x it holds that $|\{y \in [2^\ell - 1] : \mathbf{nmE}(x, y) = 1\}| > \epsilon' \cdot 2^\ell$. Further suppose that $B : [2^\ell - 1] \times [2^{\ell+2} - 1] \rightarrow \{0, 1\}$ is quasi-orthogonal with error $0.1 \cdot \epsilon'$. Then, the $(6 \cdot 2^\ell - 3)$ -vertex graph $G = (V_0 \cup V_1 \cup V_2, E)$ such that $V_\sigma = \{\langle \sigma, i \rangle : i \in [2^{\ell_\sigma} - 1]\}$, where $\ell_0 = \ell_1 = \ell$ and $\ell_2 = \ell + 2$, and

$$E = \{\{\langle 1, i \rangle, \langle 0, j \rangle\} : \mathbf{nmE}(i, j) = 1\} \cup \{\{\langle \sigma, i \rangle, \langle 2, j \rangle\} : B(i, j) = 1, \sigma \in \{0, 1\}\} \cup \binom{V_1}{2} \cup \binom{V_2}{2} \quad (22)$$

is $\Omega(|V|)$ -robustly self-ordered, where $V = V_0 \cup V_1 \cup V_2$. Furthermore, each vertex in this graph has degree at least $0.3 \cdot |V|$ and at most $0.9 \cdot |V|$. Moreover, the claim holds even if the non-malleability condition (i.e., Eq. (18)) holds only for permutations f and g that have no fixed-points.

Using the foregoing ingredients (including the non-malleable extractor of [7, Thm. 1]), Theorem 1.4 follows (see also Remark 8.9). Looking at Eq. (22), note that the first set of edges corresponds to a bipartite graph between V_1 and V_0 that is determined by \mathbf{nmE} , the second set corresponds to the bipartite graphs between V_σ (for $\sigma \in \{0, 1\}$) and V_2 that are determined by B , and the other two sets correspond to cliques on V_1 and on V_2 (see Figure 2).

Proof: Recall that $V = V_0 \cup V_1 \cup V_2$, and consider an arbitrary (non-trivial) permutation $\mu : V \rightarrow V$. Intuitively, if μ maps a vertex of V_0 (or V_1) to V_2 , then the difference in the degrees of the vertices in the two sets (caused by the $|V_2|$ -clique edges) contributes $\Omega(|V|)$ units to the symmetric difference between G and $\mu(G)$, where here we use the first quasi-orthogonality condition of B . A similar

argument, which uses the V_1 -clique edges and relies on the linear degrees of \mathbf{nmE} , applies to a vertex of V_σ mapped to $V_{1-\sigma}$ for any $\sigma \in \{0, 1\}$. On the other hand, if for some $\sigma \in \{0, 1, 2\}$ the bijection μ maps $\langle \sigma, i \rangle \in V_\sigma$ to $V_\sigma \setminus \{\langle \sigma, i \rangle\}$, then the difference in the neighborhoods caused by one of the two relevant bipartite graphs contributes $\Omega(|V|)$ units to the symmetric difference between G and $\mu(G)$. Here, we distinguish between the case that μ has relatively few non-fixed-points in either V_0 or V_1 , which is analyzed using the second quasi-orthogonality condition of B , and the case that μ has relatively many non-fixed-points in both V_0 and V_1 , which is analyzed using the non-malleability condition of \mathbf{nmE} . Indeed, the structure of the proof is similar to the one of Theorem 8.3, but the details are different in many aspects, and so we provide them below.

Let $T = \{v \in V : \mu(v) \neq v\}$ denote the set of non-fixed-points of μ . Then, we consider two types of vertices: Those that belong to the set $T' = \bigcup_{\sigma \in \{0,1,2\}} \{v \in V_\sigma : \mu(v) \notin V_\sigma\} \subseteq T$ and those that belong to $T \setminus T'$. The threshold for distinguishing these cases is set to $K = (0.5 - 0.1 \cdot \epsilon') \cdot |V_0|/4 = \Omega(|V|)$.⁴⁴ Recall that ϵ denotes the extraction error of \mathbf{nmE} , whereas ϵ' is the fractional degree bound associated with its linear degrees feature, and $0.1 \cdot \epsilon'$ is the quasi-orthogonality error of B .

Case 1: $|T'| \geq K$.

(This refers to the case that many vertices are mapped by μ to a different part of the three-way partition (V_0, V_1, V_2) of V , where ‘many’ means $\Omega(|V|)$.)

Each vertex in T' contributes $\Omega(|V|)$ units to the symmetric difference between G and $\mu(G)$, because of the differences in the degrees of vertices in the three parts. Specifically:

- Vertices in V_2 have degree at least $(|V_2| - 1) + (0.5 - 0.1\epsilon') \cdot (|V_0| + |V_1|) > (5 - 0.2\epsilon') \cdot |V_0| - O(1)$, where the first term is due to the clique edges and the second term is due to the bipartite graphs connecting V_2 to V_0 and to V_1 (and relies on the first quasi-orthogonality condition of B).
- Vertices in V_0 have degree at most $|V_1| + (0.5 + 0.1\epsilon') \cdot |V_2| < (3 + 0.4\epsilon') \cdot |V_0| + O(1)$, where the first term is due to the edges (determined by \mathbf{nmE}) connecting V_0 to V_1 and the second term is due to the bipartite graph connecting V_0 to V_2 .
- Vertices in V_1 have degree at least $(|V_1| - 1) + \epsilon' \cdot |V_0| + (0.5 - 0.1\epsilon') \cdot |V_2| > (3 + 0.6\epsilon') \cdot |V_0| - O(1)$ and at most $(|V_1| - 1) + |V_0| + (0.5 + 0.1\epsilon') \cdot |V_2| < (4 + 0.4\epsilon') \cdot |V_0|$. In both cases, the first term is due to clique edges, the second term is due to the edges connecting V_1 to V_0 (as determined by \mathbf{nmE}), and the third term is due to the edges connecting V_1 to V_2 (as determined by B). The crucial fact is that the linear degrees of \mathbf{nmE} provides a non-trivial lower bound (of $\epsilon' \cdot |V_0|$) on the second term.

Hence, the difference in the degrees of vertices in the different parts is at least $0.2\epsilon' \cdot |V_0| - O(1)$, where the minimum is due to the difference between the degrees of vertices in V_1 and the degrees of vertices in V_0 .

It follows that the symmetric difference between G and $\mu(G)$ is at least $(0.2\epsilon' \cdot |V_0| - O(1)) \cdot |T'| = \Omega(|V|) \cdot |T'|$, since $|V_0| = \Omega(|V|)$ and $\epsilon' = \Omega(1)$. Using the case’s hypothesis, we have $|T'| = \Omega(|V|) = \Omega(|T|)$, which means that in this case the difference between G and $\mu(G)$ is $\Omega(|V|) \cdot |T|$.

⁴⁴The threshold is set depending on the quasi-orthogonality error of B . In the proof of Theorem 8.3, the threshold was set depending on the quasi-orthogonality error of \mathbf{nmE} (which equaled its extraction error).

We stress that the difference between G and $\mu(G)$ is at least $\Omega(|V|) \cdot |T'|$ also if the case hypothesis does not hold.

Case 2: $|T'| < K$.

(This refers to the case that few vertices are mapped by μ to a different part of the three-way partition (V_0, V_1, V_2) of V .)

For every $\sigma \in \{0, 1, 2\}$, let $V'_\sigma = V_\sigma \cap \mu^{-1}(V_\sigma)$ and $T_\sigma = V'_\sigma \cap T$. Indeed, (T', T_0, T_1, T_2) is a four-way partition of T . Note that the size of the symmetric difference between G and $\mu(G)$ is lower-bounded by

$$\begin{aligned} & |\{(v, u) \in V'_1 \times V'_0 : \text{nmE}(\mu(v), \mu(u)) \neq \text{nmE}(v, u)\}| \\ & + |\{(v, u) \in V'_1 \times V'_2 : B(\mu(v), \mu(u)) \neq B(v, u)\}| \\ & + |\{(v, u) \in V'_0 \times V'_2 : B(\mu(v), \mu(u)) \neq B(v, u)\}|, \end{aligned} \quad (23)$$

since, for any $(v, u) \in V'_1 \times V'_0$, it holds that $\mu(v)$ neighbors $\mu(u)$ in G if and only if $\text{nmE}(\mu(v), \mu(u)) = 1$, whereas $\mu(v)$ neighbors $\mu(u)$ in $\mu(G)$ if and only if v neighbors u in G which holds if and only if $\text{nmE}(v, u) = 1$. Ditto for the other two cases.

We consider two sub-cases according to whether or not $\min(|T_0|, |T_1|)$ is relatively large. The threshold for distinguishing these sub-cases is also set to $K = (0.5 - 0.1 \cdot \epsilon') \cdot |V_0|/4$; note that $K = \Omega(|V|)$ and $K > 0.1 \cdot |V_0| > 2^{\ell-4} \geq 2^k$.

Case 2.1: $\min(|T_0|, |T_1|) < K$.

In this case we shall use the quasi-orthogonality of B .

Suppose, without loss of generality, that $|T_0| \leq |T_1|$, which implies $|T_0| < K$.

Depending on the relative sizes of T_1 and T_2 , we shall use either the quasi-orthogonal bipartite graph between V_1 and V_2 or the quasi-orthogonal bipartite graph between V_2 and V_0 .

1. If $|T_1| > |T_2|$, then we consider the quasi-orthogonal bipartite graph between V_1 and V_2 . The contribution of each vertex $v \in T_1$ to Eq. (23) equals

$$\begin{aligned} & |\{u \in V'_2 : B(\mu(v), \mu(u)) \neq B(v, u)\}| \\ & \geq |\{u \in V'_2 : B(\mu(v), u) \neq B(v, u)\}| - |T_2| \\ & > |\{u \in V_2 : B(\mu(v), u) \neq B(v, u)\}| - |T'| - |T_1| \\ & \geq (0.5 - 0.1 \cdot \epsilon') \cdot |V_2| - K - |V_0| \\ & \geq (4 - 0.25) \cdot (0.5 - 0.1\epsilon') \cdot |V_0| - |V_0| \\ & > 0.6 \cdot |V_0| \end{aligned}$$

where the first inequality uses $\mu(u) = u$ for $u \in V'_2 \setminus T_2$, the second inequality uses $|V'_2| \geq |V_2| - |T'|$ and the hypothesis $|T_2| < |T_1|$, the third inequality uses $\mu(v) \neq v$ along with the (second condition of) quasi-orthogonality of B (and the hypotheses $|T'| < K$ and the fact that $|T_1| \leq |V_1| = |V_0|$), the fourth inequality uses $|V_2| > 4 \cdot |V_0|$ and the definition of K , and the last inequality uses $\epsilon' < 0.5$.

So the total contribution in this sub-case is $|T_1| \cdot \Omega(|V|) \geq (|T| - |T'|) \cdot \Omega(|V|)$, since $|T_1| \geq \max(|T_0|, |T_2|)$ and $|T_0| + |T_1| + |T_2| = |T| - |T'|$.

2. If $|T_1| \leq |T_2|$, then we consider the quasi-orthogonal bipartite graph between V_2 and V_0 . The contribution of each vertex $v \in T_2$ to Eq. (23) equals

$$\begin{aligned}
& |\{u \in V'_0 : B(\mu(u), \mu(v)) \neq B(u, v)\}| \\
& \geq |\{u \in V'_0 : B(u, \mu(v)) \neq B(u, v)\}| - |T_0| \\
& \geq |\{u \in V_0 : B(u, \mu(v)) \neq B(u, v)\}| - |T'| - |T_0| \\
& \geq (0.5 - 0.1 \cdot \epsilon') \cdot |V_0| - 2 \cdot K \\
& = (0.5 - 0.1 \cdot \epsilon') \cdot |V_0|/2
\end{aligned}$$

where the first inequality uses $\mu(u) = u$ for $u \in V'_0 \setminus T_0$, the second inequality uses $|V'_0| \geq |V_0| - |T'|$, the third inequality uses $\mu(v) \neq v$ along with the (second condition of) quasi-orthogonality of B (and the hypotheses regarding $|T'|$ and $|T_0|$), and the equality is due to $K = (0.5 - 0.1 \cdot \epsilon') \cdot |V_0|/4$.

So the total contribution in this sub-case is $|T_2| \cdot \Omega(|V|) \geq (|T| - |T'|) \cdot \Omega(|V|)$, since $|T_2| \geq |T_1| \geq |T_0|$.

Hence, the total contribution (of Case 2.1) to Eq. (23) is $\Omega(|V|) \cdot (|T| - |T'|)$.

Case 2.2: $\min(|T_0|, |T_1|) \geq K$.

In this case we shall use the non-malleable feature of \mathbf{nmE} .

Specifically, for each $\sigma \in \{0, 1\}$, let μ_σ denote the restriction of μ to T_σ . Essentially, using $K \geq 2^k$, the non-malleability condition of the (k, ϵ) -extractor \mathbf{nmE} implies

$$|\{(i, j) \in T_0 \times T_1 : \mathbf{nmE}(i, j) \neq \mathbf{nmE}(\mu_0(i), \mu_1(j))\}| \geq (0.5 - \epsilon) \cdot |T_0| \cdot |T_1|.$$

As in (Case 2.2 of) the proof of Theorem 8.3, this can be seen by letting X and Y be uniform over T_0 and T_1 , respectively.⁴⁵

Hence, in this case, the total contribution to Eq. (23) is $(0.5 - \epsilon) \cdot |T_0| \cdot |T_1| = \Omega(|V|^2)$, where we use $\min(|T_0|, |T_1|) = \Omega(|V|)$.

Hence, in both sub-cases, the difference between G and $\mu(G)$ is $\Omega(|V|) \cdot (|T| - |T'|)$.

Recall that (by the last comment at Case 1) the difference between G and $\mu(G)$ is $\Omega(|V|) \cdot |T'|$. Combining this lower-bound with the conclusion of Case 2, the difference between G and $\mu(G)$ is $\Omega(|V|) \cdot |T|$. As for the degree bounds, note that each vertex has degree at most $(|V_2| - 1) + (0.5 + 0.1\epsilon') \cdot (|V_0| + |V_1|) = \frac{5+0.2\epsilon'}{6} \cdot |V| + O(1)$, and at least $(0.5 - 0.1\epsilon') \cdot |V_2| = \frac{2-0.4\epsilon'}{6} \cdot |V| - O(1)$, where maximum (resp., minimum) is obtained by vertices in V_2 (resp., V_0). ■

Digest: Compared to the construction used in Theorem 8.3, the construction in Theorem 8.7 decouples the non-malleable feature from the quasi-orthogonality feature, using non-malleable extractors for connecting one pair of vertex-sets and quasi-orthogonal functions to connect the other two pairs. The current analysis is slightly more complex because it has to handle the fact that these features hold for different pairs. Specifically, the quasi-orthogonality of B is used in Cases 1

⁴⁵Recall that, formally, we also extend μ_0 and μ_1 (which have no fixed-points) to (arbitrary) derangements f and g , respectively. Then, the non-malleability condition implies that $(\mathbf{nmE}(X, Y), \mathbf{nmE}(f(X), g(Y)))$ is ϵ -close to $(U_1, \mathbf{nmE}(f(X), g(Y)))$, which implies that $\Pr[\mathbf{nmE}(X, Y) \neq \mathbf{nmE}(\mu_0(X), \mu_1(Y))]$ is at least $\Pr[U_1 \neq \mathbf{nmE}(\mu_0(X), \mu_1(Y))] - \epsilon = 0.5 - \epsilon$.

and 2.1, whereas the non-malleability of `nmE` is used in Case 2.2. In particular, Case 1 only uses the first condition of quasi-orthogonality, and does so in order to infer that the degrees of all vertices in the bipartite graph determined by B are approximately equal. In Case 2.1 the second quasi-orthogonality condition is used in order to assert that the neighborhoods of two different vertices in V_σ (for every $\sigma \in \{0, 1, 2\}$) are significantly different. This is useful only when the number of non-fixed-points in the other side of the graph B is relatively small.

In light of the key role that quasi-orthogonal unbalanced bipartite graphs play in Theorem 8.7 and given their natural appeal, it feels adequate to provide a general construction of these graphs, which generalizes the construction outlined before Theorem 8.7 (for the case of $\ell' = \ell + 2$).

Proposition 8.8 (quasi-orthogonal unbalanced bipartite graphs): *For $\ell \leq \ell'$ and $S_\ell \stackrel{\text{def}}{=} \{0, 1\}^\ell \setminus \{0^\ell\}$ let $G : S_\ell \rightarrow S_{\ell'}$ be a small-bias generator with bias ϵ such that $G(s) \neq G(s')$ for every $s \neq s'$, and let E_2 denote the inner-product mod 2 function. Then, the bipartite graph described by the adjacency predicate $B : S_\ell \times S_{\ell'} \rightarrow \{0, 1\}$ such that $B(x, y) = E_2(G(x), y)$ is quasi-orthogonal with error ϵ .*

(Note that the hypothesis implies that $\epsilon > 1/|S_{\ell'}|$. The definition of quasi-orthogonal bipartite graphs appears before Theorem 8.7.)

Proof Sketch: Our starting point is the fact that $E_2 : S_{\ell'} \times S_{\ell'} \rightarrow \{0, 1\}$ is quasi-orthogonal with error $1/|S_{\ell'}|$. The quasi-orthogonality feature of the first argument of B follows as a special case of the corresponding feature of E_2 . Turning to fixings of the second argument of E_2 and letting X be uniform over S_ℓ , we observe that, for every $y \in S_{\ell'}$, the bit $B(G(X), y)$ is a linear combination of the bits of $G(X)$, and hence $\Pr[B(G(X), y) = 1] = 0.5 \pm \epsilon$. Similarly, for $y \neq y'$, it holds that $B(G(X), y) \oplus B(G(X), y') = B(G(X), y \oplus y')$ is linear combination of the bits of $G(X)$. ■

Remark 8.9 (obtaining $\Omega(n)$ -robustly self-ordered n -vertex graphs, for every n): *Theorem 8.7 provides a construction of $\Omega(n)$ -robustly self-ordered n -vertex graphs, for every n of the form $6 \cdot 2^\ell - 3$, where $\ell \in \mathbb{N}$. A construction for every $n \in \mathbb{N}$ can be obtained by using a few minor modifications.*

- *Rather than using $|V_2| = 2^{\ell+2} - 1 = 4 \cdot (|V_0| + 1) - 3$, we may use $|V_2| = n - 2 \cdot |V_0|$ such that $|V_0| = \Omega(n)$. Specifically, we still use $|V_0| = 2^\ell - 1$, for $\ell = \log_2 n - \Theta(1)$, along with $|V_2| \in [4 \cdot |V_0|, 10 \cdot |V_0|]$. Doing so requires decreasing the quasi-orthogonality error of B to $0.04\epsilon'$ so that $0.04\epsilon' \cdot |V_2| \leq 0.4 \cdot |V_0|$ still holds.*
- *More importantly, we need a construction of a quasi-orthogonal bipartite graph with an adjacency predicate $B : [2^\ell - 1] \times [n'] \rightarrow \{0, 1\}$ such that $n' = n - 2 \cdot (2^\ell - 1) \geq 2n/3$. The solution is to associate $[n']$ with an easily enumerable small-bias space $S \subseteq \{0, 1\}^{\ell+4} \setminus \{0^{\ell+4}\}$ and use $B(x, y) = E_2(G(x), y)$, where E_2 and G are as in Proposition 8.8. Specifically, for $t = \log_2 \log_2 \ell$ and $D = \lceil n' \cdot 2^t / 2^{\ell+4} \rceil = \Theta(2^t)$, we let S contain the n' lexicographically first strings in $S' \times \{0, 1\}^{\ell+4-t}$, where S' is a small-bias sample space of size D over $\{0, 1\}^t$ that is found by exhaustive search.⁴⁶*

⁴⁶Note that for every $z = (z', z'') \in \{0, 1\}^{\ell+4} \setminus \{0^{\ell+4}\}$ and $Y = (Y', Y'')$ that is uniformly distributed over S such that $|z'| = |Y'| = t$ it holds that

$$\mathbb{E}[(-1)^{E_2(z, Y)}] = \mathbb{E}[(-1)^{E_2(z', Y')}] \cdot \mathbb{E}[(-1)^{E_2(z'', Y'')}]$$

where the absolute value of each of the factors is $o(1)$ if the corresponding fixed string (i.e., z' or z'') is non-zero. Specifically, note that Y' (resp., Y'') is $o(1)$ -close to being uniformly distributed over S' (resp., $\{0, 1\}^{\ell+4-t}$).

8.3 Obtaining efficient self-ordering

We say that a self-ordered graph $G = ([n], E)$ is **efficiently self-ordered** if there exists a polynomial-time algorithm that, given any graph $G' = (V', E')$ that is isomorphic to G , finds the unique bijection $\phi : V' \rightarrow [n]$ such that $\phi(G') = G$ (i.e., the unique isomorphism of G' and G). Indeed, this isomorphism orders the vertices of G' in accordance with the original (or target) graph G .

Recall that in the case of bounded-degree graphs, we relied on the existence of a polynomial-time isomorphism test (see [30]) for efficiently self-ordering the robustly self-ordered graphs that we constructed. We cannot do so in the dense graph case, since a general polynomial-time isomorphism test is not known (see [1]). Instead, we augment the construction asserted in Theorem 1.4 so to obtain dense $\Omega(n)$ -robustly self-ordered graphs that are efficiently self-ordered.⁴⁷

Theorem 8.10 (strengthening Theorem 1.4): *There exist an infinite family of dense $\Omega(n)$ -robustly self-ordered graphs $\{G_n\}_{n \in \mathbb{N}}$ and a polynomial-time algorithm that, given $n \in \mathbb{N}$ and a pair of vertices $u, v \in [n]$ in the n -vertex graph G_n , determines whether or not u is adjacent to v in G_n . Furthermore, these graphs are efficiently self-ordered, and the degrees of vertices in G_n reside in $[0.06n, 0.73n]$.*

Proof: Our starting point is the construction of m -vertex graphs that are $\Omega(m)$ -robustly self-ordered (see Theorem 1.4, which uses Theorem 8.7). Recall that the vertices in these graphs have degree that ranges between $0.3 \cdot m$ and $0.9 \cdot m$ (see Theorem 8.7).

The idea is to use two such graphs, G_1 and G_2 , one with m vertices and the other with $4 \cdot m$ vertices, where $m = n/5$, and connect them in a way that assists finding the ordering of vertices in each of these two graphs. Specifically, we designate a set, denoted S_1 , of $s \stackrel{\text{def}}{=} 2\sqrt{\log_2 n}$ vertices in $G_1 = ([m], E_1)$, and a set, denoted S_2 , of $\ell \stackrel{\text{def}}{=} \binom{s}{2} \in [1.5 \cdot \log_2 n, 2 \cdot \log_2 n]$ vertices in $G_2 = (\{m+1, \dots, 5m\}, E_2)$, and use them as follows (see Figure 3):

- Connect each vertex in S_2 to two different vertices in S_1 , while noting that each vertex in S_1 is connected to $2\ell/s = o(\ell)$ vertices of S_2 .
- Connect each vertex in $R_1 \stackrel{\text{def}}{=} [m] \setminus S_1$ to a different set of neighbors in S_2 such that each vertex in R_1 has at least $\ell/2$ neighbors in S_2 . (Note that $\binom{\ell}{\ell/2} > n$.)
- Connect each vertex in $R_2 \stackrel{\text{def}}{=} \{m+1, \dots, 5m\} \setminus S_2$ to a different set of neighbors in R_1 such that each vertex in R_2 has two neighbors in R_1 and each vertex in R_1 has at most eight neighbors in R_2 .

Denote the resulting graph by $G = ([n], E)$, and note that the vertices of G_1 have degree at most $0.9 \cdot m + \ell$, whereas the vertices of G_2 have degree at least $0.3 \cdot 4m$. Given an isomorphic copy of the G , we can find the unique isomorphism (i.e., its ordering) as follows:

⁴⁷Unlike in the bounded degree case (see Section 4.4), at the time of writing this paper, we did not know how to construct $\Omega(n)$ -robustly self-ordered graphs that support *local* self-ordering, but such constructions were presented in subsequent works. Specifically, $\Omega(n)$ -robustly self-ordered graphs with information-theoretically local self-ordering were presented in our follow-up work [24], and a construction that supports (full-fledged) local self-ordering was presented by the first author [18]. We mention that the latter work provides a general study of the relationship between robust self-ordering and local self-ordering (in both regimes).

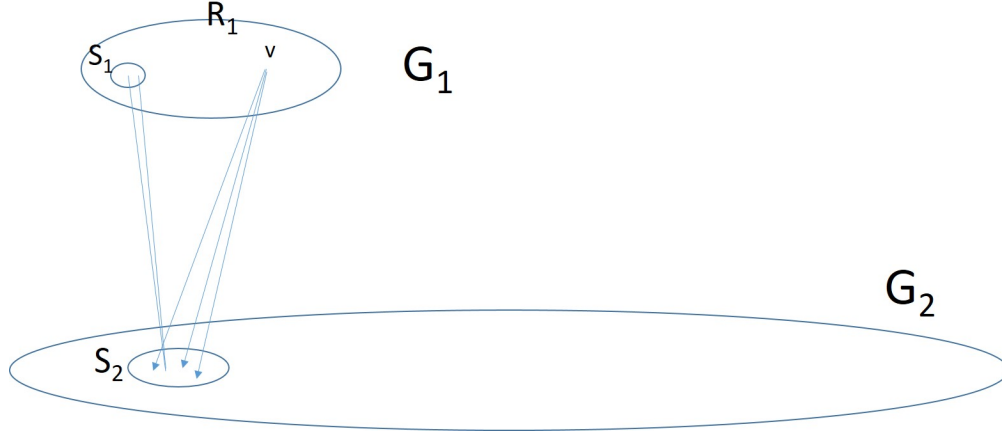


Figure 3: The construction of Theorem 8.10.

1. Identify the vertices that belong to G_1 by virtue of their lower degree.
2. Identify the set S_1 as the set of vertices that belong to G_1 and have $2\ell/s = o(\ell)$ neighbors in G_2 .
(Recall that each vertex in R_1 has at least $\ell/2$ neighbors in S_2 .)
3. Identify the set S_2 as the set of vertices that belong to G_2 and have (two) neighbors in S_1 .
4. For each possible ordering of S_1 , order the vertices of S_2 by their neighborhood in S_1 , and order the vertices of R_1 according to their neighborhood in S_2 .
(Recall that the index of each $s \in S_2$ is determined by its (two) neighbors in S_1 , whereas the index of each $v \in R_1$ is determined by its neighbors in S_2 .)
If the resulting ordering (of $S_1 \cup R_1$) yields an isomorphism to G_1 , then continue. Otherwise, try the next ordering of S_1 .
5. Order the vertices of R_2 according to their neighborhood in R_1 .

Note that by the asymmetry of G_1 , there exists a unique ordering of its vertices, and a unique ordering of S_1 that fits it and leads the procedure to successful termination. On the other hand, the number of possible ordering of S_1 is $s! = n^{o(1)}$, which means that the procedure is efficient.

It is left to show that the graph G is $\Omega(n)$ -robustly self-ordered. Let $\gamma \in (0, 1]$ be a constant such that that G_1 (resp., G_2) is $\gamma \cdot m$ -robustly self-ordered (resp., $\gamma \cdot 4m$ -robustly self-ordered). Then, fixing an arbitrary permutation $\mu : [n] \rightarrow [n]$, and letting $T = \{v \in [n] : \mu(v) \neq v\}$, we consider the following cases.

Case 1: $t \stackrel{\text{def}}{=} |\{v \in [m] : \mu(v) \notin [m]\}| > \gamma \cdot |T|/10$.

In this case, we get a contribution of $t \cdot \Omega(m) = \Omega(m \cdot |T|)$ units to the symmetric difference between G and $\mu(G)$, because of the difference in degree between vertices in $[m]$ and outside $[m]$. (Recall that the former have degree at most $0.9 \cdot m + \ell < m$, whereas the latter have degree at least $0.3 \cdot 4m = 1.2 \cdot m$.)

Case 2: $t = |\{v \in [m] : \mu(v) \notin [m]\}| \leq \gamma \cdot |T|/10$.

In this case, at least $(1 - 0.1\gamma) \cdot |T|$ vertices in T are mapped by μ to the side in which they belong (i.e., each of these vertices v satisfies $v \in [m]$ if and only if $\mu(v) \in [m]$). Let $T_1 \stackrel{\text{def}}{=} \{v \in T \cap [m] : \mu(v) \in [m]\}$ and $T_2 \stackrel{\text{def}}{=} \{v \in T \setminus [m] : \mu(v) \notin [m]\}$. Then, the vertices in T_1 contribute at least $|T_1| \cdot \gamma \cdot m - t \cdot m$ units to the symmetric difference between G and $\mu(G)$, where the negative term is due to possible change in the incidence with vertices that did not maintain their side. Similarly, the vertices in T_2 contribute at least $|T_2| \cdot \gamma \cdot 4m - t \cdot 4m$ units to the symmetric difference. Hence, in total, we get a contribution of at least $(|T| - 2t) \cdot \gamma \cdot m - t \cdot 5m = \Omega(m \cdot |T|)$.

The main claims follow. As for the degree bounds, note that the degree of each vertex in G_1 is at least $0.3m = 0.06n$, whereas the degree of each vertex in G_2 is at most $0.9 \cdot 4m + s < 0.73n$. ■

Digest. The n -vertex graph constructed in the proof of Theorem 8.10 is proved to be $\Omega(n)$ -robustly self-ordered by implicitly using the following claim.

Claim 8.11 (combining two $\Omega(n)$ -robustly self-ordered graphs): *For $i \in \{1, 2\}$, let $G_i = (V_i, E_i)$ be an $\Omega(n)$ -robustly self-ordered graph, and consider a graph $G = (V_1 \cup V_2, E_1 \cup E_2 \cup E)$ such that E contain edges with a single vertex in each V_i ; that is, G consists of G_1 and G_2 and an arbitrary bipartite graph that connects them. If the maximum degree in G of each vertex in V_1 is smaller by an $\Omega(n)$ term from the minimum degree of each vertex in V_2 , then G is $\Omega(n)$ -robustly self-ordered.*

Indeed, Claim 8.11 is analogous to Claim 4.3 (which refers to bounded-degree graphs).

We also comment that $\Omega(n)$ -robustly self-ordered graph maintain this feature also when $o(n)$ edges are added (and/or removed) from the incidence of each vertex. A related statement appears as Theorem 10.4.

9 Application to Testing Dense Graph Properties

In Section 5, we demonstrated the applicability of robustly self-ordered bounded-degree graphs to the study of testing graph properties in the bounded-degree graph model. In the current section, we provide an analogous demonstration for the regime of dense graphs. Hence, we refer to testing graph properties in the dense graph model, which was introduced in [19] and is surveyed in [16, Chap. 8]. In this model, graphs are represented by their adjacency predicate, and distances are measured as the ratio of the number of differing adjacencies to the maximal number of edges.

Background. We represent a graph $G = ([n], E)$, by the adjacency predicate $g : [n] \times [n] \rightarrow \{0, 1\}$ such that $g(u, v) = 1$ if and only if $\{u, v\} \in E$, and oracle access to a graph means oracle access to its adjacency predicate (equiv., adjacency matrix). The distance between the graphs $G = ([n], E)$ and $G' = ([n], E')$ is defined as the fraction of entries (in the adjacency matrix) on which the two graphs disagree.

Definition 9.1 (testing graph properties in the dense graph model): *A tester for a graph property Π is a probabilistic oracle machine that, on input parameters n and ϵ , and oracle access to an n -vertex graph $G = ([n], E)$ outputs a binary verdict that satisfies the following two conditions.*

1. If $G \in \Pi$, then the tester accepts with probability at least $2/3$.
2. If G is ϵ -far from Π , then the tester accepts with probability at most $1/3$, where G is ϵ -far from Π if for every n -vertex graph $G' = ([n], E') \in \Pi$ the adjacency matrices of G and G' disagree on more than $\epsilon \cdot n^2$ entries.

The query complexity of a tester for Π is a function (of the parameters n and ϵ) that represents the number of queries made by the tester on the worst-case n -vertex graph, when given the proximity parameter ϵ .

Our result. We present a general reduction of testing any property Φ of (bit) strings to testing a corresponding graph property Π . Loosely speaking, n -bit long strings will be encoded as part of an $O(\sqrt{n})$ -vertex graph, which is constructed using $\Omega(\sqrt{n})$ -robustly self-ordered $\Theta(\sqrt{n})$ -vertex graphs. This reduction is described in Construction 9.2 and its validity is proved in Lemma 9.3. Denoting the query complexities of Φ and Π by Q_Φ and Q_Π , respectively, we get $Q_\Phi(n, \epsilon) \leq Q_\Pi(O(n^{1/2}), \Omega(\epsilon))$. Thus, lower bounds on the query complexity of testing Φ , which is a property of “ordered objects” (i.e., bit strings), imply lower bounds on the query complexity of testing Π , which is a property of “unordered objects” (i.e., graphs).

Our starting point is the construction of m -vertex graphs that are $\Omega(m)$ -robustly self-ordered. Actually, wishing Π to preserve the computational complexity of Φ , we use a construction of graphs that are efficiently self-ordered, as provided by Theorem 8.10. Recall that the vertices in these graphs have degree that ranges between $0.06 \cdot m$ and $0.73 \cdot m$.

The idea is to use two such graphs, G_1 and G_2 , one with m vertices and the other with $49 \cdot m$ vertices, where $m = \sqrt{n}$, and encode an n -bit string in the connection between them. Specifically, we view the latter string as a m -by- m matrix, denoted $(s_{i,j})_{i,j \in [m]}$, and connect the i^{th} vertex of G_1 to the j^{th} vertex of G_2 if and only if $s_{i,j} = 1$.

Construction 9.2 (from properties of strings to properties of dense graphs): *Suppose that $\{G_m = ([m], E_m)\}_{m \in \mathbb{N}}$ is a family of $\Omega(m)$ -robustly self-ordered m -vertex graphs. For every $n \in \mathbb{N}$, we let $m = \sqrt{n}$, and proceed as follows.*

- For every $s \in \{0, 1\}^n$ views as $(s_{i,j})_{i,j \in [m]} \in \{0, 1\}^{m \times m}$, we define the graph $G'_s = ([50m], E'_s)$ such that

$$E'_s = E_m \cup \{\{m+i, m+j\} : \{i, j\} \in E_{49m}\} \cup \{\{i, m+j\} : i, j \in [m] \wedge s_{i,j} = 1\} \quad (24)$$

That is, G'_s consists of a copy of G_m and a copy of G_{49m} that are connected by a bipartite graph that is determined by s .

- For a set of strings Φ , we define $\Pi = \bigcup_{n \in \mathbb{N}} \Pi_n$ as the set of all graphs that are isomorphic to some graph G'_s such that $s \in \Phi$; that is,

$$\Pi_n = \{\pi(G'_s) : s \in (\Phi \cap \{0, 1\}^n) \wedge \pi \in \text{Sym}_{50m}\} \quad (25)$$

where Sym_{50m} denote the set of all permutations over $[50m]$.

Note that, given a graph of the form $\pi(G'_s)$, the vertices of G_m are easily identifiable (as having degree at most $0.73m + m < 1.8m$).⁴⁸ The foregoing construction yields a local reduction of Φ to Π , where locality means that each query to G'_s can be answered by making a constant number of queries to s . The (standard) validity of the reduction (i.e., $s \in \Phi$ if and only if $G'_s \in \Pi$) is based on the fact that G_m and G_{49m} are asymmetric.

In order to be useful towards proving lower bounds on the query complexity of testing Π , we need to show that the foregoing reduction is “distance preserving” (i.e., strings that are far from Φ are transformed into graphs that are far from Π). The hypothesis that G_m and G_{49m} are $\Omega(m)$ -robustly self-ordered is pivotal to showing that if the string s is far from Φ , then the graph G'_s is far from Π .

Lemma 9.3 (preserving distances): *If $s \in \{0, 1\}^n$ is ϵ -far from Φ , then the $50m$ -vertex graph G'_s (as defined in Construction 9.2) is $\Omega(\epsilon)$ -far from Π .*

Proof: We prove the contrapositive. Suppose that G'_s is δ -close to Π . Then, for some $r \in \Phi$ and a permutation $\pi : [50m] \rightarrow [50m]$, it holds that G'_s is δ -close to $\pi(G'_r)$, which means that these two graphs differ on at most $\delta \cdot (50m)^2$ vertex pairs. If $\pi(i) = i$ for every $i \in [2m]$, then s must be $O(\delta)$ -close to r , since $s_{i,j} = 1$ (resp., $r_{i,j} = 1$) if and only if i is connected to $m + j$ in G'_s (resp., in $\pi(G'_r) = G'_r$).⁴⁹ Unfortunately, the foregoing condition (i.e., $\pi(i) = i$ for every $i \in [2m]$) need not hold in general.

In general, the hypothesis that $\pi(G'_r)$ is δ -close to G'_s implies that π maps at most $O(\delta m)$ vertices of $[m]$ to $\{m + 1, \dots, 2m\}$, and maps to $[m]$ at most $O(\delta m)$ vertices that are outside it. This is the case because each vertex of $[m]$ has degree smaller than $0.73m + m < 1.8m$, whereas the other vertices have degree at least $0.06 \cdot 49m > 2.9m$.

Turning to the vertices $i \in [m]$ that π maps to $[m] \setminus \{i\}$, we upper-bound their number by $O(\delta m)$, since the difference between $\pi(G'_r)$ and G'_s is at most $\delta \cdot (50m)^2$, whereas the hypothesis that G_m is $c \cdot m$ -robustly self-ordered implies that the difference between $\pi(G'_r)$ and G'_s (or any other graph G'_w) is at least

$$\Delta = c \cdot m \cdot |\{i \in [m] : \pi(i) \neq i\}| - m \cdot |\{i \in [m] : \pi(i) \notin [n]\}|.$$

(Hence, $|\{i \in [m] : \pi(i) \neq i\}| \leq \frac{\Delta + m \cdot O(\delta m)}{cm} = O(\delta m)$.) The same considerations apply to the vertices $i \in \{m + 1, \dots, 2m\}$ that π maps to $\{m + 1, \dots, 2m\} \setminus \{i\}$; their number is also upper-bounded by $O(\delta m)$.

For every $k \in \{1, 2\}$, letting $I_k = \{i \in [m] : \pi((k - 1) \cdot m + i) = (k - 1) \cdot m + i\}$, observe that $D \stackrel{\text{def}}{=} |\{(i, j) \in I_0 \times I_1 : r_{i,j} \neq s_{i,j}\}| \leq \delta \cdot (50m)^2$, since $r_{i,j} \neq s_{i,j}$ implies that $\pi(G'_r)$ and G'_s differ on the vertex-pair $(i, m + j)$. Recalling that $m - |I_k| = O(\delta m)$, it follows that

$$|\{(i, j) \in [m] : r_{i,j} \neq s_{i,j}\}| \leq ((m - |I_1|) - (m - |I_2|)) \cdot m + D = O(\delta m^2).$$

Hence, s is $O(\delta)$ -close to $r \in \Phi$, and the claims follows. \blacksquare

⁴⁸In contrast, the vertices of G_{49m} have degree at least $0.06 \cdot 49m > 2.9m$.

⁴⁹Hence, G'_s is δ -close to G'_r implies that $|\{i, j \in [n] : s_{i,j} \neq r_{i,j}\}| \leq \delta \cdot (50m)^2$, which means that s is $\frac{(50m)^2 \delta}{n}$ -close to r . (Recall that $m = \sqrt{n}$.)

10 The Case of Intermediate Degree Bounds

While Section 2–6 study bounded-degree graphs and Sections 7–9 study dense graphs (i.e., constant edge density), in this section we shall consider graphs of intermediate degree bounds. That is, for every $d : \mathbb{N} \rightarrow \mathbb{N}$ such that $d(n) \in [\Omega(1), n]$, we consider n -vertex graphs of degree bound $d(n)$. In this case, the best robustness we can hope for is $\Omega(d(n))$, and we shall actually achieve it for all functions d .

Theorem 10.1 (robustly self-ordered graphs for intermediate degree bounds): *For every $d : \mathbb{N} \rightarrow \mathbb{N}$ such that $d(n)$ is computable in $\text{poly}(n)$ -time, there exists an efficiently constructable family of graphs $\{G_n\}_{n \in \mathbb{N}}$ such that G_n has maximal degree $d(n)$ and is $\Omega(d(n))$ -robustly self-ordered.*

We prove Theorem 10.1 in three parts, each covering a different range of degree-bounds (i.e., $d(n)$'s). Most of the range (i.e., $d(n) = \Omega(\log n)^{0.5}$) is covered by Theorem 10.2, whereas Theorem 10.3 handles small degree-bounds (i.e., $d(n) = O(\log n)^{0.499}$) and Theorem 10.5 handles the degree-bounds that are in-between. One ingredient in the proof of Theorem 10.5 is a transformation of graphs that makes them expanding, while preserving their degree and robustness parameters up to a constant factor. This transformation, which is a special case of Theorem 10.4, is of independent interest.

Theorem 10.2 (robustly self-ordered graphs for large degree bounds): *For every $d : \mathbb{N} \rightarrow \mathbb{N}$ such that $d(n) \geq O(\sqrt{\log n})$ is computable in $\text{poly}(n)$ -time, there exists an efficiently constructable family of graphs $\{G_n\}_{n \in \mathbb{N}}$ such that G_n has maximal degree $0.79 \cdot d(n)$ and is $\Omega(d(n))$ -robustly self-ordered. Furthermore, the minimal vertex degree is $0.02 \cdot (n)$.*

The graphs will consist of connected components of size $d(n)$, and in this case $d(n) = \Omega(\sqrt{\log n})$ is necessary, since these connected components must be different.

Proof Sketch: We combine ideas from Construction 9.2 with elements of the proof of Theorem 4.2. Specifically, as in Construction 9.2, we shall use constructions of m -vertex and $9m$ -vertex graphs that are $\Omega(m)$ -robustly self-ordered, but here we set $m = d(n)/10$ (rather than $m = \sqrt{n}$) and use $n/d(n)$ different $d(n)$ -vertex graphs that are based on the foregoing two graphs. As in the proof of Theorem 4.2, these ($d(n)$ -vertex) graphs will be far from being isomorphic to one another and will form the connected components of the final n -vertex graph.

Our starting point is the construction of m -vertex graphs that are $\Omega(m)$ -robustly self-ordered. Specifically, we may use Theorem 8.6 and note that in this case the vertices in these m -vertex graph have degree that ranges between $0.24 \cdot m$ and $0.76 \cdot m$. Furthermore, these graphs have extremely high conductance; that is, in each of these graphs, the number of edges crossing each cut (in the graph) is at least $\Omega(m)$ times the number of vertices in the smaller side (of the cut).

The idea is to use two such graphs, G_1 and G_2 , one with $m \stackrel{\text{def}}{=} 0.1 \cdot d(n)$ vertices and the other with $0.9 \cdot d(n) = 9 \cdot m$ vertices, and connect them in various ways as done in Section 4.2. Specifically, using an error correcting code with constant rate and constant relative distance and weight, denoted $C : [2^k] \rightarrow \{0, 1\}^{m^2}$, we obtain a collection of $2^k \geq n/d(n)$ strongly connected $d(n)$ -vertex graphs such that the i^{th} graph consists of copies of G_1 and G_2 that are connected according to the codeword $C(i)$; more specifically, we use the codeword $C(i)$ (viewed as an m -by- m matrix) in order to determine the connections between the vertices of G_1 and the first $0.1 \cdot d(n)$

vertices of G_2 . The final n -vertex graph, denoted G , consists of $n/d(n)$ connected components that are the first $n/d(n)$ graphs in this collection.⁵⁰

The analysis adapts the analysis of the construction presented in the proof of Theorem 4.2. Towards this analysis, we let $G_j^{(i)}$ denote the i^{th} copy of G_j ; that is, the copy of G_j that is part of the i^{th} connected component of G . Hence, for each $i \in [n/d(n)]$, the i^{th} connected component of G is isomorphic to a graph that consists of copies of $G_1 = ([m], E_1)$ and $G_2 = (\{m+1, \dots, 10m\}, E_2)$ such that for every $u, v \in [m]$ the vertex u (of $G_1^{(i)}$) is connected to the vertex $m+v$ (of $G_2^{(i)}$) if and only if $C(i)_{u,v} = 1$. Loosely speaking, considering an arbitrary permutation $\mu : [n] \rightarrow [n]$, we proceed as follows.⁵¹

- The discrepancy between the degrees of vertices in copies of G_1 and G_2 (i.e., degree smaller than $0.76m + m$ versus degree at least $0.24 \cdot 9m$) implies that each vertex that resides in a copy of G_1 and is mapped by μ to a copy of G_2 yields a contribution of $\Omega(d(n))$ units to the symmetric difference between G and $\mu(G)$.
- Let $\mu'(i)$ (resp., $\mu''(i)$) denote the index of the connected component to which μ maps a plurality of the vertices that reside in $G_1^{(i)}$ (resp., of $G_2^{(i)}$). Then, the extremely high conductance of G_1 (resp., G_2) implies that the vertices that reside in $G_1^{(i)}$ (resp., of $G_2^{(i)}$) and are mapped by μ to a connected component different from $\mu'(i)$ (resp., $\mu''(i)$) yields an average contribution of $\Omega(d(n))$ units per each of these vertices.
- The lower bound on the weight of the codewords of C , which implies that there are $\Omega(d(n)^2)$ edges between $G_1^{(i)}$ and $G_2^{(i)}$, implies that every i such that $\mu'(i) \neq \mu''(i)$ yields a contribution of $\Omega(d(n)^2)$ units, since there are no edges between $G_1^{(\mu'(i))}$ and $G_2^{(\mu''(i))}$. Here we assume that few vertices fell to the previous case (i.e., are mapped by μ in disagreement with the relevant plurality vote); analogously to the proof of Theorem 4.2, each of these few exceptional vertices reduces the contribution by at most $d(n)$ units.
- The $\Omega(d(n))$ -robust self-ordering of G_1 (resp., G_2) implies that each vertex that reside in $G_1^{(i)}$ (resp., of $G_2^{(i)}$) and is mapped by μ to a different location in $G_1^{(\mu'(i))}$ (resp., in $G_2^{(\mu''(i))}$) yields a contribution of $\Omega(d(n))$ units. Again, this assumes that few vertices fell to the penultimate case, whereas each of these few vertices reduces the contribution by one unit (per each vertex in the current case).
- The distance between the codewords of C implies that every i such that $\mu'(i) = \mu''(i) \neq i$ yields a contribution of $\Omega(d(n)^2)$, where we assume that few vertices fell to the previous cases.

As in the proof of Theorem 4.2, there may be a double counting across the different cases, but this only means that we overestimate the contribution by a constant factor. Overall the size of the symmetric difference is $\Omega(d(n))$ times the number of non-fixed-points of μ . ■

⁵⁰Note that we used $2^k \geq n/d(n)$ and $m^2 = O(k)$, where $m = 0.1 \cdot d(n) > \sqrt{k}$. This setting allows for handling any $d(n) \geq O(\sqrt{\log n})$.

⁵¹These cases are analogous to the cases treated in the proof of Theorem 4.2, with the difference that we merged Cases 2&3 (resp., Cases 4&5) into our second (resp., third) case.

Handling small degree bounds. Theorem 10.2 is applicable only for degree bounds that are at least $O(\log n)^{0.5}$. A different construction allows handling degree bounds up to $O(\log n)^{0.499}$, which leaves a small gap (which we shall close in Theorem 10.5).

Theorem 10.3 (robustly self-ordered graphs for small degree bounds): *For every constant $\epsilon > 0$, and every $d : \mathbb{N} \rightarrow \mathbb{N}$ such that $d(n) \in [\Omega(1), (\log n)^{0.5-\epsilon}]$ is computable in $\text{poly}(n)$ -time, there exists an efficiently constructable family of graphs $\{G_n\}_{n \in \mathbb{N}}$ such that G_n has maximal degree $d(n)$ and is $\Omega(d(n))$ -robustly self-ordered.*

In this case, the graphs will consist of connected components of size $\frac{\Theta(\log n)}{d(n) \cdot \log \log n} > d(n)$.

Proof Sketch: Setting $m(n) \stackrel{\text{def}}{=} \frac{\Theta(\log n)}{d(n) \cdot \log \log n} > d(n) \cdot (\log n)^\epsilon$, we proceed in two main steps.

The first main step amounts to showing that, with probability at least $1 - \exp(-\Omega(d(n) \cdot \log m(n))) = 1 - o(1)$, a $d(n)$ -regular $m(n)$ -vertex multi-graph (generated by the random permutation model) is $\Omega(d(n))$ -robustly self-ordered and expanding. This is a tightening of Theorem 6.1 (where $m(n) = n$ and $d(n) = O(1)$) and it is proved by observing that the original proof (as is) extends to a varying degree bound and that higher robustness can be obtained as indicated in Footnote 34. Specifically, we can show that each non-fixed-point i contributes at least $d(n)/2$ units (rather than one) to the symmetric difference, by considering smaller sets of indices $j \in [2d(n)]$ for which equalities hold (i.e., we use $|J_i| \geq d(n)/12$ instead of $|J_i| \geq d(n)/6$). The second main step consists of mimicking the construction outlined in Remark 6.2; specifically, we use the following three steps.

1. First, we extend the argument of the foregoing main step and show that, for any set \mathcal{G} of $t < n$ multi-graphs (which is each $d(n)$ -regular and has $m(n)$ vertices), with probability at least $1 - \exp(-\Omega(d(n) \cdot \log m(n)) - t \cdot \exp(\Omega(d(n) \cdot m(n) \cdot \log(m(n)/d(n)))) = 1 - o(1)$, a random $d(n)$ -regular $m(n)$ -vertex multi-graph (as generated above) is both $\Omega(d(n))$ -robustly self-ordered (and expanding) and far from being isomorphic to any multi-graph in \mathcal{G} . Here two $d(n)$ -regular $m(n)$ -vertex multi-graphs are said to be **far apart** if they disagree on $\Omega(d(n) \cdot m(n))$ vertex-pairs. As in Step 1 of Remark 6.2, this amounts to observing that the probability that such a random multi-graph is close to being isomorphic to a fixed multi-graph is at most $\exp(-\Omega(d(n) \cdot m(n) \cdot \log(m(n)/d(n)))) = o(1/n^2)$, where the last inequality is due to the setting of $m(n)$.⁵²

(Note that this multi-graph may have parallel edges and self-loops, but their number can be upper-bounded with high probability. Specifically, for $t = 1/\epsilon$, with probability at least $1 - O(d(n)^t/m(n)^{t-1})$, no vertex has t (or more) self-loops and no vertex is incident to $t + 1$ (or more) parallel edges. Hence, omitting all self-loops and all parallel edges leaves us with a simple graph that is both $\Omega(d(n))$ -robustly self-ordered (and expanding) and far from being isomorphic to any graph in \mathcal{G} .)

2. Next, using Step 1, we show that one can construct in $\text{poly}(n)$ -time a collection of $n/m(n)$ graphs such that each graph is $d(n)$ -regular, has $m(n)$ vertices, is $\Omega(d(n))$ -robustly self-ordered and expanding, and the graphs are pairwise far from being isomorphic to one another.

⁵²For starters, the probability that an edge that appears in the fixed multi-graph appears in the random graph is $d(n)/m(n)$. Intuitively, these events are sufficiently independent so to prove the claim; for example, we may consider the neighborhoods of the first $m(n)/2$ vertices in the random graph, and use an iterative process in which their incidences are determined at random conditioned on all prior choices.

As in Step 2 of Remark 6.2, this is done by iteratively finding $\Omega(d(n))$ -robustly self-ordered $d(n)$ -regular $m(n)$ -vertex expanding graphs that are far from being isomorphic to all prior ones, while relying on the fact that $m(n)^{d(n) \cdot m(n)} = \text{poly}(n)$ (by the setting of $m(n)$).

3. Lastly, we use the graphs constructed in Step 2 as connected components of an n -vertex graph, and obtain the desired graph.

Note that we have used $m(n) > (\log n)^\epsilon \cdot d(n)$ and $d(n) \cdot m(n) \cdot \log(m(n)/d(n)) = \Theta(\log n)$, which is possible if (and only if) $d(n) \leq (\log n)^{0.5 - \Theta(\epsilon)}$. ■

Obtaining strongly connected graphs. The graphs constructed in the proofs of Theorems 10.2 and 10.3 consists of many small connected components; specifically, we obtain n -vertex graphs of maximum degree $d(n)$ with connected components of size $\max(O(d(n)), o(\log n))$ that are $\Omega(d(n))$ -robustly self-ordered. We point out that the latter graphs can be transformed into ones with asymptotically maximal expansion (under any reasonable definition of this term), while preserving their maximal degree and robustness parameter (up to a constant factor). This is a consequence of the following general transformation (applied to a $\Omega(d(n))$ -robustly self-ordered graph G of maximal degree $d(n)$ and a $d'(n)$ -regular graph G' of maximal expansion).

Theorem 10.4 (the effect of super-imposing two graphs): *For every $d, d' : \mathbb{N} \rightarrow \mathbb{N}$ and $\rho : \mathbb{N} \rightarrow \mathbb{R}$, let G and G' be n -vertex graphs such that G is $\rho(n)$ -robustly self-ordered and has maximum degree $d(n)$, and G' has maximum degree $d'(n)$. Then, the graph obtained by super-imposing G and G' is $(\rho(n) - d'(n))$ -robustly self-ordered and has maximum degree $d(n) + d'(n)$.*

Note that Theorem 10.4 is not applicable to the constructions of bounded-degree graphs obtained in the first part of this paper, because their robustness parameter was a constant smaller than 1. (This is due mostly to Construction 2.3, but also occurs in the proof of Theorem 4.2.)⁵³ A typical application of Theorem 10.4 may use $d'(n) = \rho(n)/2 \geq 3$. (Recall that $\rho(n) \leq d(n)$ always holds.)

Proof: Fixing any permutation μ of the vertex set, note that the contribution of each non-fixed-point of μ to the symmetric difference between $G \cup G'$ and $\mu(G \cup G')$ may decrease by at most $d'(n)$ units due to G' . ■

Closing the gap between Theorems 10.2 and 10.3. Recall that these theorems left few bounding functions untreated; essentially, these were functions $d : \mathbb{N} \rightarrow \mathbb{N}$ such that $d(n) \in [(\log n)^{0.499}, O(\log n)^{0.5}]$. We close this gap now.

Theorem 10.5 (robustly self-ordered graphs for the remaining degree bounds): *For every $d : \mathbb{N} \rightarrow \mathbb{N}$ such that $d(n) \in [(\log n)^{1/3}, (\log n)^{2/3}]$ is computable in $\text{poly}(n)$ -time, there exists an efficiently constructable family of graphs $\{G_n\}_{n \in \mathbb{N}}$ such that G_n has maximal degree $d(n)$ and is $\Omega(d(n))$ -robustly self-ordered.*

In this case, the graphs will consist of connected components of size $O(\log n)$.

⁵³In contrast, the construction of Theorem 10.3, which builds upon the proof of Theorem 6.1, does yield $\Omega(d)$ -robustly self-ordered graphs of maximum degree d , for sufficiently large constant d .

Proof Sketch: We apply the proof strategy of Theorem 10.2, while setting $\ell = \log_2 n$ and using as building blocks ℓ -vertex $\Omega(d(n))$ -robustly self-ordered graphs of degree bound $d(n)/2$ obtained from Theorem 10.2 itself (while relying on $d(n) \in [\omega(\log \ell)^{1/2}, o(\ell)]$). Actually, we shall use two versions of the foregoing ℓ -vertex graphs that have sufficiently different degree ranges and transform them to obtain maximal expansion (using Theorem 10.4). Furthermore, we shall connect pairs of copies of these graphs by subgraphs of a fixed $d(n)/2$ -regular bipartite graph (rather than by subgraphs of the complete ℓ -by- ℓ bipartite graph). As in the proof of Theorem 10.2, using sufficiently different bipartite subgraphs, we shall obtain 2ℓ -vertex graphs that are far from being isomorphic to one another. Our final n -vertex graph will consist of these 2ℓ -vertex graphs as its connected components. Details follow.

Our starting point is a construction of ℓ -vertex graphs that, for some constant $\gamma \in (0, 0.01)$, are $\gamma \cdot d(n)$ -robustly self-ordered and have maximum degree $0.79 \cdot d(n)$ and minimum degree $0.02 \cdot d(n)$. Such graphs are obtained by Theorem 10.2, while relying on $d(n) \geq \ell^{1/3} = \omega(\log \ell)$ and $d(n) < \ell$. Using Theorem 10.4 (with $d'(n) = \gamma \cdot d(n)/2$) we transform these graphs to ones of maximum degree $0.8d(n)$ and asymptotically maximal conductance (i.e., in each of these graphs, the number of edges crossing each cut (in the graph) is at least $\Omega(d(n))$ times the number of vertices in the smaller side (of the cut)). Note that the resulting graph, denoted G_1 , is $\Omega(d(n))$ -robustly self-ordered and has minimum degree $0.02 \cdot d(n)$. Using the same process with degree bound $d_2(n) = 0.01 \cdot d(n)$, we obtain an analogous graph, denoted G_2 , that is also $\Omega(d(n))$ -robustly self-ordered and has maximum degree $0.8 \cdot d_2(n) < 0.01 \cdot d(n)$. We stress that both G_1 and G_2 have asymptotically maximal conductance (w.r.t degree bound $\Theta(d(n))$).

Next, we connect G_1 and G_2 in various ways so to obtain $n/2\ell$ graphs that are far from being isomorphic to one another. This is done by a small variation on the construction used in the proof of Theorem 10.2. Specifically, we fix $d(n)/200$ disjoint perfect matchings between the vertices of G_1 and the vertices G_2 , and use the error correcting code to determine which of the corresponding $\ell \cdot d(n)/200 = \omega(\log n)$ edges to include in the graph. More specifically, using an error correcting code with constant rate and constant relative distance and weight, denoted $C : [2^k] \rightarrow \{0, 1\}^{\ell \cdot d(n)/200}$, we obtain a collection of $n/2\ell < 2^k$ strongly connected 2ℓ -vertex graphs such that the i^{th} graph consists of copies of G_1 and G_2 that are connected according to the codeword $C(i)$; that is, the $(j, v)^{\text{th}}$ bit of the codeword $C(i)$ (viewed as an $d(n)/200$ -by- ℓ matrix) determines whether the edge of the j^{th} matching that is incident at vertex $v \in [\ell]$ (of G_1) is included in the i^{th} graph. The final n -vertex graph, denoted G , consists of these $n/2\ell$ graphs as its connected components.

The analysis is almost identical to the analysis provided in the proof of Theorem 10.2, since the key facts used there hold here too (although the construction is somewhat different). The key facts are that the degrees of vertices in G_1 and G_2 differ in $\Omega(d(n))$ units, that the relative conductance of the connected components is $\Omega(d(n))$, that G_1 and G_2 are both $\Omega(d(n))$ -robustly self-ordered, and that the bipartite graphs (used in the different connected components) are far away from one another. ■

Acknowledgements

We are grateful to Eshan Chattopadhyay for discussions regarding non-malleable two-source extractors, and to Dana Ron for discussions regarding tolerant testing. We are deeply indebted to Jyun-Jie Liao for pointing out an error in the proof of [22, Lem. 9.7]. We are also grateful to the reviewers of this journal for their extremely careful reading of the paper and for their numerous comments.

Oded Goldreich is a Meyer W. Weisgal Professor at the Faculty of Mathematics and Computer Science of the Weizmann Institute of Science, Rehovot, ISRAEL. His research was partially supported by the Israel Science Foundation (grant No. 1041/18) and by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No. 819702). Avi Wigderson is the Herbert Maass Professor at the School of Mathematics of the Institute for Advanced Study in Princeton, NJ 08540, USA. His research was partially supported by NSF grant CCF-1900460.

References

- [1] L. Babai. Graph Isomorphism in Quasipolynomial Time [extended abstract]. In *48th ACM Symposium on the Theory of Computing*, pages 684–697, 2016.
- [2] E. Ben-Sasson, O. Goldreich, P. Harsha, M. Sudan, and S. Vadhan. Robust PCPs of Proximity, Shorter PCPs, and Applications to Coding. *SIAM Journal on Computing*, Vol. 36 (4), pages 889–974, 2006.
- [3] A. Bogdanov, K. Obata, and L. Trevisan. A Lower Bound for Testing 3-Colorability in Bounded-Degree Graphs. In *43rd IEEE Symposium on Foundations of Computer Science*, pages 93–102, 2002.
- [4] B. Bollobas. Distinguishing Vertices of Random Graphs. *North-Holland Mathematics Studies*, Vol. 62, pages 33–49, 1982.
- [5] B. Bollobas. The Asymptotic Number of Unlabelled Regular Graphs. *J. Lond. Math. Soc.*, Vol. 26, pages 201–206, 1982.
- [6] J. Bourgain and A. Gamburd. Uniform expansion bounds for Cayley graphs of $SL_2(F_p)$. *Annals of Mathematics*, pages 625–642, 2008.
- [7] E. Chattopadhyay, V. Goyal, and X. Li. Non-Malleable Extractors and Codes, with their Many Tampered Extensions. In *48th STOC*, pages 285–298, 2016.
- [8] M. Cheraghchi and V. Guruswami. Non-Malleable Coding Against Bit-Wise and Split-State Tampering. In *11th TCC*, pages 440–464, 2014.
- [9] B. Chor and O. Goldreich. Unbiased Bits from Sources of Weak Randomness and Probabilistic Communication Complexity. *SIAM Journal on Computing*, Vol. 17, No. 2, pages 230–261, 1988.
- [10] I. Dinur, O. Goldreich, and T. Gur. Every Set in P Is Strongly Testable Under a Suitable Encoding. In *10th ITCS*, pages 30:1–30:17, 2019.

- [11] I. Dinur and O. Reingold. Assignment-testers: Towards a combinatorial proof of the PCP-Theorem. *SIAM Journal on Computing*, Vol. 36 (4), pages 975–1024, 2006.
- [12] Y. Dodis and D. Wichs. Non-malleable extractors and symmetric key cryptography from weak secrets. In *41st STOC*, pages 601–610, 2009.
- [13] D. Ellis. Lecture 13: The Expansion of Random Regular Graphs. Lecture notes, Algebraic Methods in Combinatorics, University of Cambridge, March 2011. Available at <https://davidellis2.files.wordpress.com/2019/07/lecture13.pdf>
- [14] P. Erdos and A. Renyi. Asymmetric Graphs. *Acta Mathematica Hungarica*, Vol. 14 (3), pages 295–315, 1963.
- [15] E. Fischer and L. Fortnow. Tolerant Versus Intolerant Testing for Boolean Properties. *Theory of Computing*, Vol. 2 (9), pages 173–183, 2006.
- [16] O. Goldreich. *Introduction to Property Testing*. Cambridge University Press, 2017.
- [17] O. Goldreich. On Testing Hamiltonicity in the Bounded Degree Graph Model. *ECCC*, TR20-109, 2020.
- [18] O. Goldreich. Robust Self-Ordering versus Local Self-Ordering. *ECCC*, TR21-034, 2021.
- [19] O. Goldreich, S. Goldwasser, and D. Ron. Property testing and its connection to learning and approximation. *Journal of the ACM*, pages 653–750, July 1998. Extended abstract in *37th FOCS*, 1996.
- [20] O. Goldreich, M. Krivelevich, I. Newman, and E. Rozenberg. Hierarchy Theorems for Property Testing. *Computational Complexity*, Vol. 21 (1), pages 129–192, 2012.
- [21] O. Goldreich and D. Ron. Property Testing in Bounded Degree Graphs. *Algorithmica*, Vol. 32 (2), pages 302–343, 2002.
- [22] O. Goldreich and A. Wigderson. Robustly Self-Ordered Graphs: Constructions and Applications to Property Testing. *ECCC*, TR20-149, September 2020.
- [23] O. Goldreich and A. Wigderson. Constructing Large Families of Pairwise Far Permutations: Good Permutation Codes Based on the Shuffle-Exchange Network. *ECCC*, TR20-192, 2020.
- [24] O. Goldreich and A. Wigderson. Non-adaptive vs Adaptive Queries in the Dense Graph Testing Model. *ECCC*, TR20-160, 2020.
- [25] C.S. Greenhill, S. Janson, J.H. Kim, and N.C. Wormald. Permutation Pseudographs and Contiguity. *Combinatorics, Probability and Computing*, Vol. 11, pages 273–298, 2002.
- [26] S. Hoory, N. Linial, and A. Wigderson. Expander Graphs and Their Applications. Bulletin (New Series) of the American Mathematical Society, Vol. 43 (4), pages 439–561, 2006.
- [27] J.H. Kim, B. Sudakov, and V.H. Vu. On the asymmetry of random regular graphs and random graphs. *Random Structures & Algorithms*, Vol. 21 (3-4), pages 216–224, 2002.

- [28] A. Lubotzky. *Discrete Groups, Expanding Graphs and Invariant Measures*. Progress in mathematics, Vol. 125, Birkhauser, 1994.
- [29] A. Lubotzky, R. Phillips, and P. Sarnak. Ramanujan Graphs. *Combinatorica*, Vol. 8, pages 261–277, 1988.
- [30] E.M. Luks. Isomorphism of Graphs of Bounded Valence can be Tested in Polynomial Time. *Journal of Computer and System Science*, Vol. 25 (1), pages 42–65, 1982.
- [31] M. Parnas, D. Ron, and R. Rubinfeld. Tolerant property testing and distance approximation. *Journal of Computer and System Science*, Vol. 72(6), pages 1012–1042, 2006.
- [32] R. Rubinfeld and M. Sudan. Robust characterization of polynomials with applications to program testing. *SIAM Journal on Computing*, Vol. 25(2), pages 252–271, 1996.