The Share Size of Secret-Sharing Schemes for Almost All Access Structures and Graphs

Amos Beimel*
Ben-Gurion University of the Negev
Be'er-Sheva, Israel

Oriol Farràs†
Universitat Rovira i Virgili
Tarragona, Spain

June 3, 2020

Abstract

The share size of general secret-sharing schemes is poorly understood. The gap between the best known upper bound on the total share size per party of $2^{0.64n+o(n)}$ (Applebaum et al., STOC 2020) and the best known lower bound of $\Omega(n/\log n)$ (Csirmaz, J. of Cryptology 1997) is huge (where n is the number of parties in the scheme). To gain some understanding on this problem, we study the share size of secret-sharing schemes of almost all access structures, i.e., of almost all collections of authorized sets. This is motivated by the fact that in complexity, many times almost all objects are hardest (e.g., most Boolean functions require exponential size circuits). All previous constructions of secret-sharing schemes were for the worst access structures (i.e., all access structures) or for specific families of access structures.

We prove upper bounds on the share size for almost all access structures. We combine results on almost all monotone Boolean functions (Korshunov, Probl. Kibern. 1981) and a construction of (Liu and Vaikuntanathan, STOC 2018) and conclude that almost all access structures have a secret-sharing scheme with share size $2^{\tilde{0}(\sqrt{n})}$.

We also study graph secret-sharing schemes. In these schemes, the parties are vertices of a graph and a set can reconstruct the secret if and only if it contains an edge. Again, for this family there is a huge gap between the upper bounds – $O(n/\log n)$ (Erdös and Pyber, Discrete Mathematics 1997) – and the lower bounds – $\Omega(\log n)$ (van Dijk, Des. Codes Crypto. 1995). We show that for almost all graphs, the share size of each party is $n^{o(1)}$. This result is achieved by using robust 2-server conditional disclosure of secrets protocols, a new primitive introduced and constructed in (Applebaum et al., STOC 2020), and the fact that the size of the maximal independent set in a random graph is small. Finally, using robust conditional disclosure of secrets protocols, we improve the total share size for all very dense graphs.

1 Introduction

A dealer wants to store a string of secret information (a.k.a. a secret) on a set of computers such that only some pre-defined subsets of the computers can reconstruct the information. We will refer to the computers as the parties, their number as n, and the collection of authorized sets that can reconstruct the secret as

*This work was done while visiting Georgetown University, supported by NSF grant no. 1565387, TWC: Large: Collaborative: Computing Over Distributed Sensitive Data and also supported by ERC grant 742754 (project NTSC), by ISF grant 152/17, and by a grant from the Cyber Security Research Center at Ben-Gurion University of the Negev.

†This work is supported by the grant 2017 SGR 705 from the Government of Catalonia and the grant RTI2018-095094-B-C21 CONSENT from the Spanish Government.

an access structure. To achieve this goal the dealer uses a secret-sharing scheme – a randomized function that is applied to the secret and produces n strings, called shares. The dealer gives the i-th share to the i-th party, and any authorized set of parties can reconstruct the secret from its shares. Nowadays, secret-sharing schemes are used as a building box in many cryptographic tasks (in addition to their obvious usage for secure storage), e.g., secure multiparty computation protocols [21, 34, 36], threshold cryptography [44], access control [64], and attribute-based encryption [53, 73]. We consider schemes where unauthorized sets of parties gain absolutely no information on the secret from their shares, i.e., the security is information theoretic. We will mainly try to reduce the sizes of the shares given to the parties. To understand why minimizing the share size is important, let us consider the original secret-sharing schemes of [55] for an arbitrary access structure; in these schemes the size of each share is greater than 2^n, making them impractical when, for example, $n = 100$. Even in the most efficient scheme known today, the share size is $2^{0.64n}$ [5] (improving on [59, 4]).

We ask the question if the above share size can be reduced for almost all access structures. One motivation for this question is that in complexity theory, almost all Boolean functions are often the hardest functions. For example, Shannon [70] showed that almost all Boolean functions require circuits of size $2^{\Omega(n)}$, this lower bound applies also to other models, e.g., formulas. Furthermore, almost all *monotone* Boolean functions require monotone circuits and monotone formulas of size $2^{\Omega(n)}$. Dealing with properties of almost all objects is a common theme in combinatorics, e.g., properties of almost all graphs. A famous example states that the size of the maximum independent set (and clique) of almost all n-vertex graphs is approximately $2 \log n$ [54]; we use this property in our constructions. Using a result on almost all monotone Boolean functions [58], we show that almost all access structures can be realized by a secret-sharing scheme with share size $2^{\tilde{O}(\sqrt{n})}$.

In this paper, we also study graph secret-sharing schemes. In a secret-sharing scheme realizing a graph G, the parties are vertices of the graph G and a set can reconstruct the secret if and only if it contains an edge. The naive scheme to realize a graph is to share the secret independently for each edge; this result implies a share of size $O(n)$ per party. A better scheme with share size $O(n/\log n)$ per party is implied by a result of Erdös and Pyber [47]. Graph secret-sharing schemes were studied in many papers, e.g., [31, 32, 33, 45, 25, 39, 43, 40, 41, 13, 49, 42]. One motivation for studying graph secret-sharing schemes is that they are simpler than secret-sharing schemes for general access structures and phenomena proved for graph secret-sharing schemes were later generalized to general access structures (e.g., Blundo et al. [27] proved that in any non-ideal access structure the share size of at least one party is at least 1.5 times the size of the secret, a result that was later proved for every access structure [62]). Another motivation is that, by [66], for every constant $0 < c < 1/2$ any graph secret-sharing scheme with share size $O(n^c)$ per party implies a secret-sharing scheme for any access structure with share size $2^{O(0.5+c)n}$; thus, major improvement in the share size for all graphs will result in improved schemes for all access structures. However, in spite of the recent improvements in the share size for general access structures [59, 4, 5] and for specific families of access structures (e.g., forbidden graphs [17, 51, 60] and uniform access structures [2, 18, 4]), no such improvement was achieved for schemes for graphs. We show that almost all graphs can be realized by a secret-sharing scheme with share size $n^{o(1)}$ per party.

1.1 Previous Results

We next describe the most relevant previous results. We refer the reader to Figure 1 for a description of the share size in previous constructions and our constructions.

	Share size (one bit secret)	Share size of linear schemes over $\mathbb{F}_q$ ($\log q$-bit secret)	Inf. ratio multi-linear schemes (long secrets)
Forbidden graphs	$n^{o(1)}$ [60] $\Omega(1)$	$\tilde{O}(\sqrt{n}\log q)$ [51] $\Omega(\sqrt{n})$ [14]	$O(\log n)$ [2] $\Omega(1)$
Almost all graphs	$\boldsymbol{n^{o(1)}}$ Th. 5.1 $\Omega(\log\log n)$ [39]	$\boldsymbol{\tilde{O}(\sqrt{n}\log q)}$ Th. 5.1 $\Omega(\sqrt{n})$ [14]	$\boldsymbol{\tilde{O}(\log n)}$ Th. 5.1 $\boldsymbol{\Omega(\log^{1/2} n)}$ Th. 5.5
All graphs	$O(n/\log n)$ [47] $\Omega(\log n)$ [45, 40]	$O(\frac{n}{\log n}\log q)$ [47] $\Omega(\sqrt{n}\log q)$ [16]	$O(n/\log n)$ [47] $\Omega(\sqrt{n})$ [11, 16]
Almost all access structures	$\boldsymbol{2^{O(\sqrt{n\log n})}}$ Th. 3.3 $\Omega(1)$	$\boldsymbol{2^{0.5n+o(n)}\log q}$ Th. 3.3 $\boldsymbol{\Omega(2^{n/2-o(n)})}$ [9], Th. 3.10 $\boldsymbol{\Omega(2^{n/3-o(n)}\log q)}$ [68], Th. 3.9	$\boldsymbol{2^{O(\sqrt{n\log n})}}$ Th. 3.3 $\Omega(1)$
All access structures	$2^{0.64n}$ [5] $\Omega(n/\log n)$ [38]	$2^{0.76n}\log q$ [5] $2^{\Omega(n)}\log q$ [67]	$2^{0.64n}$ [5] $n^{\Omega(\log n)}$ [11, 16]

Figure 1: A summary of the upper and lower bounds on the maximum share size for secret-sharing schemes for forbidden graph access structures, almost all graph access structures, graph access structures, almost all access structures, and all access structures. The results proved in this paper are in **boldface**.

Measures of share size. The size of a share is simply the length of the string representing it. For a secret-sharing scheme, two measures of for the share size were considered: (1) the maximum share size, i.e., the maximum over all parties in the scheme of the size of the share of the party, (2) the total share size, i.e., the sum over all parties in the scheme of the size of the share of the party. For a given scheme, the maximum share size is bounded from above by the total share size, which is bounded from above by n times the maximum share size. The distinction between these two measures is important for graph secret-sharing schemes, and there might be trade-offs between optimizing one measure and optimizing the other. On the other hand, the share size in the secret-sharing schemes considered in this paper for general access structures is larger than $2^{\sqrt{n}}$, thus for these schemes the distinction between the measures is less important.

We will also consider the normalized total (respectively, maximum) share size, i.e., the ratio between the sum of the share sizes (respectively, maximum share size) and the size of the secret. This normalized share size (also known as information ratio) is similar in spirit to measures considered in information theory and it is interesting since the length of each share is at least the length of the secret [57]. In this work, we will consider the normalized share size for two regimes: (1) Moderately short secrets of size $\tilde{O}(n)$, and (2) Following [3, 2], we also consider exponentially long secrets of size 2^{n^2}. The latter size is not reasonable, however, these schemes may lead to schemes with the same share size for shorter secrets and they provide barriers for proving lower bounds via information inequalities.

Bounds on the share size. Secret-sharing schemes were introduced by Blakely [24] and Shamir [69] for the threshold case and by Ito, Saito, and Nishizeki [55] for the general case. In the original secret-sharing schemes for arbitrary access structures of Ito et al. [55] the maximum share size is 2^{n-1}. Additional constructions of secret-sharing schemes followed, e.g., [71, 30, 22, 56, 23]. For specific access structures, the share size in these schemes is less than the share size in the scheme of [55]; however, the share size in the above schemes for arbitrary access structures is $2^{n-o(n)}$. In a recent breakthrough work, Liu, and Vaikuntanathan [59] (using results of [61]) constructed a secret-sharing scheme for arbitrary access structures with

share size $2^{0.944n}$ and a linear secret-sharing scheme with share size $2^{0.999n}$. Applebaum et al. [5] (using results of [61, 4]) improved these results, constructing a secret-sharing schemes for arbitrary access structures with share size $2^{0.637n}$ and a linear secret-sharing scheme with share size $2^{0.762n}$. It is an important open problem if the share can be improved to $2^{o(n)}$ (or even smaller). Lower bounds for secret-sharing were proven in, e.g., [33, 26, 45, 38, 37]. These lower bounds are very far from the upper bounds – the best lower bound is $\Omega(n^2/\log n)$ for the normalized total share size for an explicit access structure (proven by Csirmaz [37]).

For graph secret-sharing schemes there is also a big gap between the upper bounds and lower bounds. Erdös and Pyber [47] have proved that every graph can be partitioned into complete bipartite graphs such that each vertex is contained in at most $O(n/\log n)$ complete bipartite graphs. Blundo et al. [26] observed that this implies that the normalized maximum share size of realizing every n-vertex graph is $O(n/\log n)$ (for secrets of size $\log n$). Van Dijk [45] proved a lower bound of $\Omega(\log n)$ on the normalized maximum share size of realizing an explicit n-vertex graph. Csirmaz [39] extended this lower bound to the n-vertex Boolean cube. He observed that a lower bound of $\Omega(\log n)$ on a specific graph implies a lower bound of $\Omega(\log\log n)$ for almost all graphs (as almost all n-vertex graphs contain a copy of every $\log n$-vertex graph [29]). Furthermore, Csirmaz asked if for almost every graph there is a scheme with normalized maximum share size $o(n/\log n)$. We answer this question affirmatively by showing for almost all graphs a secret-sharing scheme with maximum share size $n^{o(1)}$.

Linear secret-sharing schemes. Linear secret-sharing schemes, introduced by [30, 56], are schemes in which the random string is a vector of elements over some finite field $\mathbb{F}_q$, the domain of secrets is also $\mathbb{F}_q$, and the shares are computed as a linear map over $\mathbb{F}_q$. Many known schemes are linear, e.g., [69, 24, 22] and the schemes for graphs implied by [47]. They are equivalent to a linear-algebraic model of computation called monotone span programs [56]. Linear secret-sharing schemes are useful as they are homomorphic: given shares of two secrets s, s', each party can locally add its shares and obtain a share of $s + s'$. For many applications of secret sharing, linearity is essential, e.g., [36, 8, 74], hence, constructing linear secret-sharing schemes is important. The size of the shares in the best known linear secret-sharing scheme is $2^{0.76n}$ [5] (improving upon [59]). Pitassi and Robere [67] proved an exponential lower bound of $2^{cn}\log q$ on the share in linear secret-sharing schemes over $\mathbb{F}_q$ for an explicit access structure of (where $0 < c < 1/2$ is a constant). Babai et al. [9] proved a lower bound of $2^{n/2-o(n)}\sqrt{\log q}$ on the share in linear secret-sharing schemes over $\mathbb{F}_q$ for almost all access structures.

Multi-linear secret-sharing schemes, introduced by [23], are a generalization of linear secret-sharing schemes in which the domain of secrets is $\mathbb{F}_q^\ell$ for some integer ℓ. In [2, 5], such schemes improve the normalized maximum share size compared to the linear secret-sharing schemes constructed in those papers (i.e., the multi-linear schemes share a longer secret while using the same share size as the linear schemes). Beimel et al. [11] proved that every lower bound proved for linear secret-sharing schemes using the Gal-Pudlák criteria [50] also applies to multi-linear secret-sharing schemes. In particular, this implies that the $n^{\Omega(\log n)}$ lower bound of [9] for the normalized maximum share size for an explicit access structure and the $\Omega(\sqrt{n})$ lower bound of [16] for the normalized maximum share size for an explicit graph access structure hold also for multi-linear secret-sharing schemes. We note that it is not clear if multi-linear secret-sharing schemes can replace linear secret-sharing schemes in many applications, e.g., in the MPC protocols of [36] that are secure against general adversarial structures.

Conditional disclosure of secrets protocols. Conditional disclosure of secrets (CDS) protocols were first defined by Gertner et al. [52]. A CDS protocol for a Boolean function f involves k servers and a referee.

Each server holds a common secret s, a common random string r, and a private input x_i; using these r, s, and x_i the i-th server computes one message (without seeing any other input or message) and sends it to the referee. The referee knowing the inputs $x_1, \ldots, x_k$ and the messages should be able to compute s if and only if $f(x_1, \ldots, x_k) = 1$. CDS protocols were used in many cryptographic applications, such as symmetric private information retrieval protocols [52], attribute based encryption [51, 8, 74], priced oblivious transfer [1], and secret-sharing schemes [59, 4, 5]. Applebaum et al. [5] defined *robust* CDS protocols (see Section 2.3) and used them to construct more efficient secret-sharing schemes for arbitrary access structures. We use robust CDS protocols to construct better schemes for almost all graphs and for all very dense graphs.

The original construction of k-server CDS protocols for general functions $f : [N]^k \to \{0, 1\}$, presented in [52], has message size $O(N^k)$ (where N is the input domain size of each server). This construction is linear. Recently, better constructions of CDS protocols for general functions have been presented. Beimel et al. [17] have shown a non-linear 2-server CDS protocol with message size $O(N^{1/2})$ and Gay et al. [51] constructed a linear 2-server CDS protocol with the same message size. Then, Liu et al. [60] have designed a 2-server non-linear CDS protocol with message size $2^{O(\sqrt{\log N \log\log N})}$ and Liu et al. [61] have constructed a k-server CDS protocol with message size $2^{\tilde{O}(\sqrt{k \log N})}$. Beimel and Peter [19] and Liu et al. [61] have constructed a linear CDS protocol with message size $O(N^{(k-1)/2})$; by [19], this bound is optimal for linear CDS protocols (up to a factor of k). Applebaum and Arkis [2] (improving on Applebaum et al. [3]) have showed that there is a CDS protocol with long secrets – of size $\Theta(2^{N^k})$ – in which the message size is 4 times the secret size. Lower bounds on the message size in CDS protocols and in linear CDS protocols have been proven in [51, 3, 6, 7].

Forbidden graph access structures. In a forbidden-graph secret-sharing scheme for a graph G, introduced by Sun and Shieh [72], the parties are the vertices of the graph G and a set is authorized if it is an edge or its size is at least 3. A a forbidden-graph secret-sharing scheme for a graph G is not harder than a graph secret-sharing realizing G: Given a secret-sharing scheme realizing a graph, one can construct a forbidden-graph secret-sharing scheme for G by giving a share of the graph secret-sharing scheme and a share of a 3-out-of-n threshold secret-sharing schemes. Furthermore, forbidden graph secret-sharing schemes are closely related to 2-server CDS protocols: Beimel et al. [17] have described a transformation from a CDS protocol for a function describing the graph G to a forbidden graph secret-sharing scheme for G in which the maximum share size of the scheme is $O(\log n)$ times the message size of the CDS protocol. Furthermore, by [17, 2], if we consider secrets of size at least $O(\log^2 n)$, then there is a transformation in which the normalized maximum share size is a constant times the message size of the CDS protocol. As a result, we get that every forbidden graph G can be realized by a secret-sharing with maximum share size $n^{o(1)}$ (using the CDS protocol of [60]), by a linear secret-sharing scheme over $\mathbb{F}_q$ with maximum share size $\tilde{O}(\sqrt{n} \log q)$ for every prime power q (using the CDS protocol of [51]), and a multi-linear secret-sharing scheme with normalized maximum share size $O(1)$ for secrets of length 2^{n^2} [2]. We nearly match these bounds for graph access structures for almost all graphs.

1.2 Our Results and Techniques

We next describe the results we achieve in this paper. We again refer the reader to Figure 1 for a description of the maximum share size in previous constructions and our constructions.

Almost all access structures. We prove upper bounds on the share size for almost all access structures, namely almost all access structures have a secret-sharing scheme with share size $2^{\tilde{O}(\sqrt{n})}$, a linear secret-sharing scheme with share size $2^{n/2+o(n)}$, and a multi-linear secret-sharing scheme with maximum share size $\tilde{O}(\log n)$ for secrets of size 2^{n^2}. Our linear secret-sharing scheme for almost all access structures are optimal (up to a factor of $2^{o(n)}$) for a one-bit secret (by a lower bound of Babai et al. [9]).

The construction for almost all access structures is a simple combination of previous results. The first result, proved by Korshunov [58] in 1981, is that in almost all access structures with n parties all minimum authorized sets are of size between $n/2-1$ and $n/2+2$, i.e., all sets of size at most $n/2-2$ are unauthorized and all sets of size at least $n/2+3$ are authorized. The second result we use, proved by Liu and Vaikuntanathan [59], is that such access structures can be realized by secret-sharing schemes with share size as above. These results are presented in Section 3.

We also prove lower bounds on the normalized share size in linear secret-sharing schemes for almost all access structures. Rónyai et al. [68] proved that for every finite field $\mathbb{F}_q$ for almost all access structures the normalized share size of linear secret-sharing schemes over $\mathbb{F}_q$ realizing the access structure is at least $\Omega(2^{n/3-o(n)})$. The result of Rónyai et al. [68] does not rule-out the possibility that for every access structures there exists some finite field $\mathbb{F}_q$ (possibly with a large q) such that the access structure can be realized by a linear secret-sharing schemes over $\mathbb{F}_q$ with small normalized share size. This could be plausible since we know that there are access structures that can be realized by an efficient linear secret-sharing scheme over one field, but require large shares in any linear secret-sharing scheme over fields with a different characteristic [20, 67]. Pitassi and Robere [67] proved that there exists an explicit access structure for which this is not true, i.e., there exists a constant $c > 0$ such that in any linear secret-sharing scheme realizing it the normalized share size is 2^{cn}. In Theorem 3.9, we prove that this is not true for almost all access structures, namely, for almost every access structure the normalized share size in any linear secret-sharing scheme realizing the access structure is $\Omega(2^{n/3-o(n)})$. Our proof uses a fairly recent result on the number of representable matroids [65].

(G,t)-graph secret-sharing schemes and robust CDS. We define a hierarchy of access structures between forbidden graph access structures and graph access structures. In a (G,t)-secret-sharing scheme, every set containing an edge is authorized and, in addition, every set of size $t+1$ is authorized. In other words, the unauthorized sets are independent sets in G of size at most t. We show that (G,t)-secret-sharing schemes are equivalent to 2-server t-robust CDS protocols. As a result, using the robust CDS protocols of [5], we get efficient (G,t)-secret-sharing schemes, e.g., schemes with maximum share size $n^{o(1)}t$. These results are presented in Section 4. We note that, for an arbitrary graph G, our (G,n)-secret-sharing scheme, which is a graph secret-sharing scheme realizing G, the share size does not improve upon the scheme of [47].

Almost all graph secret-sharing schemes. We show that for almost all graphs, there exists a secret-sharing scheme with maximum share size $n^{o(1)}$, a linear secret-sharing scheme with normalized maximum share size $\tilde{O}(\sqrt{n})$ (for moderately short secrets), and a multi-linear secret-sharing scheme with normalized maximum share size $\tilde{O}(\log n)$ for exponentially long secrets. By [16, 11], there exists a graph such that in every multi-linear secret-sharing scheme realizing the graph the normalized maximum share size is $\Omega(\sqrt{n})$, thus, we get a separation for multi-linear secret-sharing schemes between the normalized maximum share size for almost all graphs and the maximum share size of the worst graph. These results are presented in Section 5.

To construct our scheme for almost all graphs, we use the fact that if the size of every independent set in a graph G is at most t, then a (G,t)-secret-sharing scheme is a graph secret-sharing scheme realizing G.

Our construction follows from the fact that for almost every graph, the size of the maximal independent set in a random graph is $O(\log n)$ [54].

We also consider the maximum share size of random n-vertex graphs drawn from the Erdös-Rényi [48] distribution $\mathscr{G}(n,p)$, that is, each pair of vertices is independently connected by an edge with probability p. For example, $\mathscr{G}(n,1/2)$ is the uniform distribution over the n-vertex graphs. On one hand, with probability nearly 1 the size of the maximum independent set in a graph drawn from $\mathscr{G}(n,p)$ is at most $O(\frac{1}{p}\log n)$, thus, using (G,t)-secret-sharing schemes with $t = O(\frac{1}{p}\log n)$, we realize a graph in $\mathscr{G}(n,p)$ with normalized maximum share size $n^{o(1)}/p$. On the other hand, with probability nearly 1 the degree of all vertices in the graph drawn from $\mathscr{G}(n,p)$ is $O(pn)$, thus, it can be realized by the trivial secret-sharing scheme with maximum share size $O(pn)$. Combining these two schemes, the hardest distribution in our construction is $\mathscr{G}(n,1/\sqrt{n})$ for which the normalized maximum share size is $\sqrt{n}$. We do not know if there is a better secret-sharing scheme for graphs drawn from $\mathscr{G}(n,1/\sqrt{n})$ or this distribution really requires shares of size $n^{\Omega(1)}$.

Dense graph secret-sharing schemes. Following [13], we study graph secret-sharing schemes for very dense graphs, i.e., graphs with at least $\binom{n}{2} - n^{1+\beta}$ edges for some constant β. For these graphs, Beimel et al. [13] have constructed a linear secret-sharing scheme with maximum share size $\tilde{O}(n^{1/2+\beta/2})$ and another linear secret-sharing scheme with total share size $\tilde{O}(n^{5/4+3\beta/4})$. We improve on the latter result and show that all very dense graphs can be realized by a secret-sharing scheme with normalized total share size of $n^{1+\beta+o(1)}$ for moderately short secrets of size $\tilde{O}(n)$. To put this result in perspective, this total share size matches (up to a factor of $n^{o(1)}$) to the total share size of the naive secret-sharing scheme for sparse graphs with $n^{1+\beta}$ edges. These schemes are presented in Section 6.

We next describe the high-level ideas of our construction realizing a graph G with at least $\binom{n}{2} - n^{1+\beta}$ edges. If every vertex in G has degree at least $n - n^{\beta}$, then the size of every independent set in G is at most $n^{\beta}+1$, and we can use a $(G, n^{\beta}+1)$-secret-sharing schemes, resulting in normalized total share size $O(n^{1+\beta+o(1)})$. While in a graph with at least $\binom{n}{2} - n^{1+\beta}$ edges the average degree is at least $n - O(n^{\beta})$, the graph can contain vertices whose degree is small. To overcome this problem, we use an idea of [13]. We consider the set of vertices A whose degree is smallest in G and execute a secret-sharing scheme realizing the graph restricted to this set (denoted G'). We choose the size of this set such that: (1) the size of the set is small, thus, the total share size in realizing G' is small, and (2) the degree of the each vertex not in A is big, thus, we can realize the graph without the edges between vertices in A by a (G,t)-secret-sharing scheme for a relatively small t. We apply the above construction iteratively to get our scheme.

Hypergraph secret-sharing schemes. A secret-sharing realizes a hypergraph H if the parties of the scheme are the vertices of H and a set of parties can reconstruct the secret if and only if it contains a hyperedge. In this work, we construct schemes for k-hypergraphs, that is, hypergraphs whose hyperedeges are all of size k. The access structures of these schemes are also called k-homogeneous. The best secret-sharing scheme for k-hypergraphs known to date is the original scheme of [55], which have maximum share size $O(\binom{n}{k-1})$.

Extending the results explained above, we show a connection between k-hypergraph secret-sharing schemes and k-server t-robust CDS protocols. For any constant k, we show that for almost every k-hypergraph there exists a secret-sharing scheme with maximum share size is $n^{o(1)}$, a linear secret-sharing scheme with normalized maximum share size $\tilde{O}(n^{(k-1)/2})$, and a multi-linear secret-sharing scheme with normalized maximum share size $\tilde{O}(\log^{k-1} n)$ for exponentially long secrets. These schemes are presented in Section 7.

Interpretation of our results. In this work we have shown that for almost all access structures there exist secret-sharing schemes that are more efficient than the *known* secret-sharing schemes for the worst access structures. Similarly, we have constructed for almost every graph G a secret-sharing schemes realizing G that are more efficient than the *known* secret-sharing schemes realizing the worst graph. One possible conclusion from this result is that in secret-sharing schemes almost all access structures might not be the hardest access structures. Another possible interpretation is that our results may be generalized to all access structures. We note that in one case we know that the former interpretation is true: there is a graph for which the normalized maximum share size for multi-linear schemes is at least $\Omega(\sqrt{n})$ (for every size of secrets) [11, 16], while we show an upper bound for almost all graphs of $\tilde{O}(\log n)$ (for long secrets).

Open problems. Can the normalized share size of almost all access structures can be improved? We do not have any non-trivial lower-bound on the normalized share size for them. Recall that an access structure is $n/2$-uniform if all sets of size less than $n/2$ are unauthorized, all sets of size greater than $n/2$ are authorized, and sets of size exactly $n/2$ can be either authorized or unauthorized. By [4] (using results of [2]), every $n/2$-uniform access structure can be realized by a scheme with normalized maximum share size $O(n^2)$ (with exponentially long secrets). Since almost all access structures somewhat resemble uniform access structures (see Theorem 3.2), one can hope that almost every access structure can be realized by a scheme with polynomial normalized share size.

Another research problem is to study the complexity of almost all functions for other primitives with information-theoretic security, for example, private simultaneous messages (PSM) protocols, MPC protocols, MPC protocols with constant number of rounds, and private information retrieval (PIR) protocols for almost all databases. For all these primitives there is a huge gap between the known upper bounds and lower bounds on the message size. Are there more efficient protocols for any of these primitives for almost all functions than the protocols for all functions?

2 Preliminaries

In the section, we present the preliminary results needed for this work. First, we define secret-sharing schemes, linear secret-sharing schemes, graph secret-sharing schemes, and homogeneous access structures. Second, we define conditional disclosure of secrets (CDS) protocols, and robust CDS protocols. We also present several CDS and robust CDS protocols from [2, 19, 60, 61] that are used in this work. Finally, we present a short introduction to random graphs and random access structures.

2.1 Secret-Sharing Schemes

We present the definition of secret-sharing scheme as given in [35, 12]. For more information about this definition and secret-sharing in general, see [10].

Definition 2.1 (Access Structures)**.** *Let $P = \{P_1, \ldots, P_n\}$ be a set of parties. A collection $\Gamma \subseteq 2^P$ is* monotone *if $B \in \Gamma$ and $B \subseteq C$ imply that $C \in \Gamma$. An* access structure *is a monotone collection $\Gamma \subseteq 2^P$ of non-empty subsets of P. Sets in Γ are called* authorized*, and sets not in Γ are called* forbidden*.*

Definition 2.2 (Secret-Sharing Schemes)**.** *A* secret-sharing scheme Π *with domain of secrets S, such that $|S| \geq 2$, is a mapping from $S \times R$, where R is some finite set called the set of random strings, to a set of n-tuples $S_1 \times S_2 \times \cdots \times S_n$, where S_j is called the* domain of shares *of P_j. A dealer distributes a secret $s \in S$ according to Π by first sampling a random string $r \in R$ with uniform distribution, computing a vector*

of shares $\Pi(s,r) = (s_1, \ldots, s_n)$*, and privately communicating each share* s_j *to party* P_j*. For a set* $A \subseteq P$*, we denote* $\Pi_A(s,r)$ *as the restriction of* $\Pi(s,r)$ *to its* A*-entries (i.e., the shares of the parties in* A*).*

A secret-sharing scheme Π *with domain of secrets* S *realizes an access structure* Γ *if the following two requirements hold:*

CORRECTNESS. *The secret* s *can be reconstructed by any authorized set of parties. That is, for any set* $B = \left\{P_{i_1}, \ldots, P_{i_{|B|}}\right\} \in \Gamma$ *there exists a* reconstruction function $\mathrm{Recon}_B : S_{i_1} \times \cdots \times S_{i_{|B|}} \to S$ *such that for every secret* $s \in S$ *and every random string* $r \in R$*,*

$$\mathrm{Recon}_B\left(\Pi_B(s,r)\right) = s.$$

PRIVACY. *Any forbidden set cannot learn anything about the secret from its shares. Formally, for any set* $T = \left\{P_{i_1}, \ldots, P_{i_{|T|}}\right\} \notin \Gamma$ *and every pair of secrets* $s, s' \in S$*, the distributions* $\Pi_T(s,r)$ *and* $\Pi_T(s',r)$ *are identical, where the distributions are over the choice of* r *from* R *at random with uniform distribution.*

Given a secret-sharing scheme Π*, define the* size *of the secret as* $\log|S|$*, the* share size *of party* P_j *as* $\log|S_j|$*, the* maximum share size *as* $\max_{1\leq j\leq n}\{\log|S_j|\}$*, and the* total share size *as* $\sum_{j=1}^{n}\log|S_j|$*.*

A secret-sharing scheme is multi-linear if the mapping that the dealer uses to generate the shares given to the parties is linear, as we formalize at the following definition.

Definition 2.3 (Multi-Linear and Linear Secret-Sharing Schemes)**.** *Let* Π *be a secret-sharing scheme with domain of secrets* S*. We say that* Π *is a* multi-linear *secret-sharing scheme over a finite field* $\mathbb{F}$ *if there are integers* $\ell_d, \ell_r, \ell_1, \ldots, \ell_n$ *such that* $S = \mathbb{F}^{\ell_d}$*,* $R = \mathbb{F}^{\ell_r}$*,* $S_1 = \mathbb{F}^{\ell_1}, \ldots, S_n = \mathbb{F}^{\ell_n}$*, and the mapping* Π *is a linear mapping over* $\mathbb{F}$ *from* $\mathbb{F}^{\ell_d+\ell_r}$ *to* $\mathbb{F}^{\ell_1+\cdots+\ell_n}$*. We say that a scheme is linear over* $\mathbb{F}$ *if* $S = \mathbb{F}$ *(i.e., when* $\ell_d = 1$*).*

Definition 2.4 (Graph secret-sharing schemes)**.** *Let* $G = (V,E)$ *be an undirected graph with* $|V| = n$*; for simplicity we assume that* $E \neq \emptyset$*. We define* Γ_G *as the access structure whose minimal authorized subsets are the edges in* G*, that is, the unauthorized sets are independent sets in the graph. A secret-sharing scheme realizing an access structure* Γ_G *is said to be a secret-sharing scheme* realizing the graph G *and is called a* graph secret-sharing schemes*.*

These schemes are one of the main topics in this work. In this paper, we study very dense graphs, graphs with at least $\binom{n}{2} - n^{1+\beta}$ edges for some $0 \leq \beta < 1$.

We also study k*-homogeneous access structures*, which are access structures whose minimal authorized subsets are of the size k. For example, graph access structures are 2-homogeneous access structures. For $k > 2$, it is convenient to define k-homogeneous access structures from hypergraphs. A *hypergraph* is a pair $H = (V,E)$ where V is a set of vertices and $E \subseteq 2^V \setminus \{\emptyset\}$ is the set of *hyperedges*. A hypergraph is k*-uniform* if $|e| = k$ for every $e \in E$. A k-uniform hypergraph is *complete* if $E = \binom{V}{k} = \{e \subseteq V : |e| = k\}$. Observe that there is a one-to-one correspondence between uniform hypergraphs and homogeneous access structures, and that complete uniform hypergraphs correspond to threshold access structures. Given a hypergraph $H = (V,E)$, we define Γ_H as the access structure whose minimal authorized sets are the hyperedges of H.

We contrast homogeneous access structures with *uniform* access structures (studied, e.g., in [72, 18, 2, 4]). A k-uniform access structures is also described by a k-uniform hyper-graph and its authorized sets are

all the hyper-edges and all sets of size at least $k+1$. Thus, k-homogeneous access structures are harder to realize as they might contain forbidden sets of size much larger than k.[1]

2.2 Conditional Disclosure of Secrets

We define k-server conditional disclosure of secrets protocols, originally defined in [52].

Definition 2.5 (Conditional Disclosure of Secrets Protocols)**.** *Let $f : X_1 \times \cdots \times X_k \to \{0,1\}$ be a k-input function. A k-*server CDS protocol *$\mathcal{P}$ for f with domain of secrets S consists of:*

1. *A finite domain of common random strings R, and k finite message domains $M_1, \ldots, M_k$,*
2. *Deterministic message computation functions $\text{ENC}_1, \ldots, \text{ENC}_k$, where $\text{ENC}_i : X_i \times S \times R \to M_i$ for every $i \in [k]$ (we also say that $\text{ENC}_i(x_i, s, r)$ is the message sent by the i-th server to the referee), and*
3. *A deterministic reconstruction function $\text{DEC} : X_1 \times \cdots \times X_k \times M_1 \times \cdots \times M_k \to \{0,1\}$.*

We denote $\text{ENC}(x, s, r) = (\text{ENC}_1(x, s, r), \ldots, \text{ENC}(x_k, s, r))$. We say that a CDS protocol $\mathcal{P}$ is a CDS protocol for a function f if the following two requirements hold:

CORRECTNESS. *For any input $(x_1, \ldots, x_k) \in X_1 \times \cdots \times X_k$ for which $f(x_1, \ldots, x_k) = 1$, every secret $s \in S$, and every common random string $r \in R$,*

$$\text{DEC}(x_1, \ldots, x_k, \text{ENC}_1(x_1, s, r), \ldots, \text{ENC}_k(x_k, s, r)) = s.$$

PRIVACY. *For any input $x = (x_1, \ldots, x_k) \in X_1 \times \cdots \times X_k$ for which $f(x_1, \ldots, x_k) = 0$ and for every pair of secrets s, s', the distributions $\text{ENC}(x, s, r)$ and $\text{ENC}(x, s', r)$ are identical, where the distributions are over the choice of r from R at random with uniform distribution.*

The message size of a CDS protocol $\mathcal{P}$ is defined as the size of largest message sent by the servers, i.e., $\max_{1 \le i \le k} \{\log |M_i|\}$.

Next, we present the properties of three CDS protocols that are used in this work. The CDS protocol presented in Theorem 2.6 has linear properties: the messages are generated from the secret and the randomness with linear mappings. Theorem 2.6 is a particular case of Theorem 6 of [2], while Theorem 2.7 is from [60].

Theorem 2.6 ([2])**.** *For any 2-input function $f : [n] \times [n] \to \{0,1\}$ there is a 2-server CDS protocol in which, for sufficiently large secrets, i.e., secrets of size 2^{n^2}, each server communicates at most 3 bits per each bit of the secret.*

Theorem 2.7 ([60])**.** *For any 2-input function $f : [n] \times [n] \to \{0,1\}$ there is a 2-server CDS protocol with message size $n^{O(\sqrt{\log \log n / \log n})} = n^{o(1)}$.*

Theorem 2.8 ([61])**.** *For any k-input functions $f : [n]^k \to \{0,1\}$ there is a k-server CDS protocol with message size $n^{O(\sqrt{k/\log n \log(k \log n)})}$.*

[1] For example, given a secret-sharing realizing the k-homogeneous access structures of a hyper-graph H, we can realize the k-uniform access structures of H by additionally sharing the secret in a $(k+1)$-out-of-k secret-sharing scheme.

2.3 Robust Conditional Disclosure of Secrets

In a recent work [5], Applebaum et al. define a stronger notion of CDS protocols that is useful for constructing secret-sharing schemes. In a k-server CDS protocol, we assume that each server sends one message to the referee. Therefore, the referee only has access to k messages. In a *robust* k-server CDS protocol, we consider the case that the referee can have access to more than one message from some servers (generated with the same common random string), and privacy is guaranteed even if an adversary sees a bounded number of messages from each server.

Definition 2.9 (Zero sets)**.** *Let $f : X_1 \times \cdots \times X_k \to \{0,1\}$ be a k-input function. We say that a set of inputs $Z \subseteq X_1 \times \cdots \times X_k$ is a* zero set *of f if $f(x) = 0$ for every $x \in Z$. For sets $Z_1, \ldots, Z_k$, we denote $\text{ENC}_i(Z_i, s, r) = (\text{ENC}_i(x_i, s, r))_{x_i \in Z_i}$, and*

$$\text{ENC}(Z_1 \times \cdots \times Z_k, s, r) = (\text{ENC}_1(Z_1, s, r), \ldots, \text{ENC}_k(Z_k, s, r)).$$

Definition 2.10 (Robust conditional disclosure of secrets (RCDS) protocols)**.** *Let $\mathcal{P}$ be a k-server CDS protocol for a k-input function $f : X_1 \times \cdots \times X_k \to \{0,1\}$ and $Z = Z_1 \times \cdots \times Z_k \subseteq X_1 \times \cdots \times X_k$ be a zero set of f. We say that $\mathcal{P}$ is robust for the set Z if for every pair of secrets $s, s' \in S$, it holds that $\text{ENC}(Z, s, r)$ and $\text{ENC}(Z, s', r)$ are identically distributed. Let $t_1, \ldots, t_k$ be integers. We say that $\mathcal{P}$ is a $(t_1, \ldots, t_k)$-robust CDS protocol if it is robust for every zero set $Z_1 \times \cdots \times Z_k$ such that $|Z_i| \leq t_i$ for every $i \in [k]$ and it is a t-robust CDS protocol if it is $(t, \ldots, t)$-robust.*

In this work we use several constructions of robust CDS protocols presented in [4], which are based on non-robust CDS protocols. Theorem 2.11 presents linear and multi-linear robust CDS protocols in which the underlying CDS protocol is from [51]. Then, Theorem 2.12 presents a generic transformation from non-robust CDS protocols to robust CDS protocols. In this transformation, if the original CDS is linear, then the resulting robust CDS is multi-linear.

Theorem 2.11 ([5, Theorem D.5])**.** *Let $f : [N] \times [N] \to \{0,1\}$ be a function. Then, for every finite field $\mathbb{F}_q$ and every integer $t \leq N/(2\log^2 N)$, there is a linear 2-server (t, N)-robust CDS protocol for f with one element secrets in which the message size is*

$$O((t\log^2 t + \sqrt{N})t \log t \log^2 N \log q).$$

Furthermore, there is p_0 such that for every prime-power $q > p_0$ there is a multi-linear 2-server (t, N)-robust CDS protocol for f over $\mathbb{F}_q$ with secrets of size $\Theta(t^2 \log q \log t \log^3 N)$ in which the normalized message size is

$$O(t\log^2 t + \sqrt{N}).$$

Theorem 2.12 ([5, Theorem E.2])**.** *Let $f : [N]^k \to \{0,1\}$ be a k-input function, for some integer $k > 1$, and $t \leq \min\{kN/2, 2^{\sqrt{N/k}}\}$ be an integer. Assume that for some integer $m \geq 1$, there is a k-server CDS protocol $\mathcal{P}$ for f with secrets of size m in which the message size is $c(N, m)$. Then, there is a k-server t-robust CDS protocol for f with secrets of size m in which the message size is*

$$O\left(c(N,m)k^{3k-1}2^k t^k \log^{2k-1} t \log^2(N)\right).$$

If $\mathcal{P}$ is a linear protocol over $\mathbb{F}_{2^m}$, then the resulting protocol is also linear. Furthermore, there is a k-server t-robust CDS protocol for f with secrets of size $\Theta(mk^2 t \log t \log^2(N))$ in which the normalized message size is

$$O\left(\frac{c(N,m)}{m} k^{3k-3} 2^k t^{k-1} \cdot \log^{2k-2} t\right).$$

2.4 Random Graphs and Access Structures

In this work, we use several results on random graphs to construct secret-sharing schemes for almost all graphs with improved share size. First, we present the Erdös-Rényi model for random graphs [48]. For an introduction to this topic see, e.g., [28].

Let $\mathscr{G}_n$ be the family of graphs with the vertex set $V = \{1, \ldots, n\}$. Given $0 < p < 1$, the model $\mathscr{G}(n, p)$ is a probability distribution over $\mathscr{G}_n$ in which each edge is chosen independently with probability p, that is, if G is a graph with m edges, then $\Pr[\{G\}] = p^m(1-p)^{\binom{n}{2}-m}$. Note that when $p = 1/2$, any two graphs are equiprobable.

We say that *almost every graph in $\mathscr{G}(n, p)$ has a certain property Q* if $\Pr[Q] \to 1$ as $n \to \infty$. For $p = 1/2$, saying that almost every graph in $\mathscr{G}(n, p)$ has a certain property Q is equivalent to saying that the number of graphs in $\mathscr{G}_n$ satisfying Q divided by $|\mathscr{G}_n|$ tends to 1 as $n \to \infty$. In this case, we will say that *almost all* graphs satisfy Q.

Analogously, we will use the same expression for any family of access structures F_n. We say that *almost all* access structures in F_n satisfy Q if the number of access structures in F_n satisfying Q divided by $|F_n|$ tends to 1 as $n \to \infty$. In particular, we study the family of homogeneous access structures and the family of all access structures.

Next, we present some properties of the maximum independent sets of graphs in $\mathscr{G}(n, p)$. Lemma 2.13 was presented by Grimmett and McDiarmid in [54]. Several subsequent results gave more accurate bound on the size of maximum independent sets, but it is enough for our purposes. In Lemma 2.14 we give bounds to the maximum independent sets in $\mathscr{G}(n, p)$ for non-constant p. In Lemma 2.15 and Lemma 2.16 we present further properties of almost all graphs. The proofs of Lemma 2.14 and Lemma 2.15 are moved to the Appendix A.

Lemma 2.13 ([54])**.** *Let $0 < p < 1$ be a constant. Then the size of a maximum independent set in almost every graph in $\mathscr{G}(n, p)$ is smaller than $2 \log n / \log(\frac{1}{1-p}) + o(\log n)$.*

As a consequence of Lemma 2.13, the size of a maximum independent set in almost every graph in $\mathscr{G}_n$ is smaller than $(2 + o(1)) \log n$.

Lemma 2.14. *The size of a maximum independent set in almost every graph in $\mathscr{G}(n, p)$ is $O(\frac{\log n}{p})$ if $1/n \leq p \leq 1/2$, and $1 + \frac{2+o(1)}{\alpha}$ if $p = 1 - n^{-\alpha}$ for some $1/\log n \leq \alpha \leq 1$.*

With a similar proof, we can also show that for every $0 \leq \beta \leq 1 - \frac{1}{\log n}$, almost all graph with $n^{1+\beta}$ edges have maximal independent sets of size at most $O(n^{1-\beta} \log n)$, and almost all graphs with $\binom{n}{2} - n^{1+\beta}$ have maximal independent sets of size at most $1 + \frac{2+o(1)}{1-\beta}$.

Lemma 2.15. *Almost all graphs in $\mathscr{G}(n, p)$ with $p = \omega(\log n/n)$ have degree at most $2pn$.*

Lemma 2.16 ([29, Theorem 1])**.** *Almost every graph with $n = \lceil r^2 2^{r/2} \rceil$ vertices contains every graph of r vertices as an induced subgraph.*

3 Secret-sharing Schemes for Almost All Access Structures

This section is dedicated to the study of general access structures. Combining results on monotone Boolean functions by Korshunov [58] and secret-sharing schemes from [59, 2], we obtain secret-sharing schemes for almost all access structures. Then, we present lower bounds on the maximum share size for almost all access structures.

3.1 Upper Bounds for Almost All Access Structures

First, we define the family of slice access structures. These access structures have a special role in the general constructions presented in [59, 4, 5]. In Theorem 3.2, we present a family of slice access structures that contains almost all access structures. It is direct consequence of the results in [58] for monotone Boolean functions (also presented in [75, Page 99]).

Definition 3.1. *Let a, b be two integers satisfying $1 \leq a < b \leq n$. We define $S_{a,b}$ as the family of access structures Γ satisfying that, for every $A \subseteq P$:*

1. *if $|A| > b$, then $A \in \Gamma$,*
2. *if $|A| < a$, then $A \notin \Gamma$.*

Theorem 3.2. *Let $\ell = \lfloor n/2 \rfloor$. Almost all access structures (i.e., monotone collections of sets) are in $S_{\ell-1,\ell+1}$ if n is even, and in $S_{\ell-1,\ell+2}$ if n is odd.*

Theorem 3.3. *Almost all access structures can be realized by the following secret-sharing schemes.*

1. *A secret-sharing scheme with share size $2^{O(\sqrt{n}\log n)}$.*
2. *A linear secret-sharing scheme with share size $2^{n/2+o(n)}$.*
3. *A multi-linear secret-sharing scheme with normalized share size $2^{O(\sqrt{n \log n})}$ for secrets of size 2^{n^2}.*

Proof. By Theorem 3.2, constructing secret-sharing schemes for access structures in $S_{\ell-1,\ell+2}$ suffices for constructing secret-sharing schemes for almost all access structures.

Assume that for every k-input function $f : [N]^k \to \{0,1\}$ and secret of size m there is a k-server CDS protocol for f in which the message size is $c(N,m)$. By [59], for every k there is a secret-sharing scheme for $\Gamma \in S_{a,b}$ with share size at most

$$c(N,m)2^{(b-a+1)n/k}O(n)\binom{n}{a}/\binom{n/k}{a/k}^k$$

for $N = \binom{n/k}{a/k}$. In our case, $a = \lfloor \frac{n}{2} \rfloor - 1$ and $b = \lfloor \frac{n}{2} \rfloor + 2$. Taking $k = \sqrt{\frac{n}{\log n}}$,

$$c(N,m)2^{4n/k}O(n)\frac{\binom{n}{n/2-1}}{\binom{n/k}{(n-2)/2k}^k} = c(N,m)2^{4\sqrt{n\log n}}O(\text{poly}(n))\left(\frac{n}{k}\right)^{\frac{k}{2}} = c(N,m)2^{O(\sqrt{n\log n})}.$$

Taking the k-server CDS protocol with message size $c(N,m) = 2^{O(\sqrt{\log N}\log\log N} \leq 2^{O(\sqrt{n}\log n)}$ from [61], we get the first secret-sharing scheme. If we take the linear k-server CDS protocol from [19, 61] with message size $O(N^{(k-1)/2}) \leq 2^{n/2+o(n)}$, we get the second secret-sharing scheme. The third secret-sharing scheme is obtained by using the k-server CDS protocol with message size $c(N,m) \leq 4m$ from [2]. □

As a consequence of this result, Hypotheses 1 and 3 in [2] are true for almost all access structures:

Hypothesis 1 (SS is short)**.** *Every access structure over n parties is realizable with small information ratio (say $2^{o(n)}$).*

Hypothesis 2 (SS is amortizable)**.** *For every access structure over n parties, and every sufficiently long secret s, there exists a secret-sharing scheme with small information ratio (e.g., sub-exponential in n).*

3.2 Almost All Access Structures Require Long Shares in Linear secret-sharing Schemes

Rónyai et al. [68] proved that for every finite field $\mathbb{F}_q$ for almost every access structure Γ the normalized share size of linear secret-sharing schemes over $\mathbb{F}_q$ realizing Γ is at least $2^{n/3-o(n)}$. We reverse the order of quantifiers and prove that for almost every access structure Γ, for every finite field $\mathbb{F}_q$ the normalized share size of linear secret-sharing schemes over $\mathbb{F}_q$ realizing Γ is at least $2^{n/3-o(n)}$.

The rest of the section is organized as follows. We start by defining monotone span program and representable matroids; these notions are used to prove the lower bounds. We then, as a warm-up, reprove a lower bound on the size of shares of linear secret-sharing over $\mathbb{F}_2$, originally proved in [9]. Thereafter, we prove our new lower bound on the normalized share size of linear secret-sharing schemes. We end this section by proving a lower bound on the size of shares in linear secret-sharing schemes for a one bit secret over all fields.

3.2.1 Definitions

A linear secret-sharing scheme with total share size m can be described by a matrix M with m rows such that the shares are computed by multiplying M by a vector whose first coordinate is the secret s and the other coordinates are random field elements. It is convenient to describe a linear secret-sharing scheme by a monotone span program, a computational model introduced by Karchmer and Wigderson [56]. The reader is referred to [10] for more background on monotone span programs and their connections to secret sharing.

Definition 3.4 (Monotone Span Program [56])**.** *A monotone span program is a triple $\mathcal{M} = (\mathbb{F}, M, \rho)$, where $\mathbb{F}$ is a field, M is an $d \times b$ matrix over $\mathbb{F}$, and $\rho : \{1, \ldots, d\} \to \{p_1, \ldots, p_n\}$ labels each row of M by a party.*[2] *The size of $\mathcal{M}$ is the number of rows of M (i.e., d). For any set $A \subseteq \{p_1, \ldots, p_n\}$, let M_A denote the sub-matrix obtained by restricting M to the rows labeled by parties in A. We say that $\mathcal{M}$ accepts B if the rows of M_B span the vector $\vec{e}_1 = (1, 0, \ldots, 0)$. We say that $\mathcal{M}$ accepts an access structure Γ if $\mathcal{M}$ accepts a set B if and only if $B \in \Gamma$.*

Theorem 3.5 ([56])**.** *There exists a linear secret-sharing scheme over $\mathbb{F}_q$ realizing an access structure Γ with secrets of size $\log q$ and total share size $d \log q$ if and only if there exists a monotone span program $\mathcal{M} = (\mathbb{F}_q, M, \rho)$ accepting the access structure Γ such that M is an $d \times d$ matrix.*

We next define representable matroids and quote the result of [65]. For our proof, we do not need the definition of matroids; we note that they are an axiomatic abstraction of linear independency.

Definition 3.6. *A matroid representable over a field $\mathbb{F}$ is a pair (A, r), where A is a finite set, called a ground set, and $r : 2^A \to \{0, 1, \ldots, |A|\}$ is a function, called a rank function, such that there are vectors $\{v_a\}_{a \in A}$ in $\mathbb{F}^{|A|}$ for which for every $B \subseteq A$*

$$r(B) = \operatorname{rank}(\{v_a\}_{a \in B}),$$

where $\operatorname{rank}(V)$ is the linear-algebraic rank of vectors, i.e., the maximum number of linearly independent vectors in V. A representable matroid is a matroid representable over some field.

Theorem 3.7 ([65])**.** *For every $d \geq 12$, there are at most $2^{d^3/4}$ representable matroids with ground set $[d]$.*

[2] For simplicity, in this paper we label a row by a party p_j rather than by a variable x_j as done in [56].

3.2.2 A Warm-up

As a warm-up, we reprove a lower bound from [9] on the share size of linear secret-sharing scheme over $\mathbb{F}_2$.

Claim 3.8 ([9])**.** *For almost every access structure Γ with n parties the share size in every linear secret-sharing scheme over $\mathbb{F}_2$ realizing Γ with a one bit secret is at least $2^{n/2-o(n)}$.*

Proof. Assume that there is a linear secret-sharing scheme over $\mathbb{F}_2$ realizing an access structure Γ with total share size d. By Theorem 3.5, there is a monotone span program $\mathcal{M} = (\mathbb{F}_2, M, \rho)$ accepting the access structure Γ where M is matrix over $\mathbb{F}_2$ of size $d \times d$.

There are at most $n^d 2^{d^2}$ monotone span programs of size d over $\mathbb{F}_2$ (as there are 2^{d^2} matrices and n ways to label each row by a party) and each monotone span program accepts a unique access structure. Thus, for $d = 0.5 \cdot 2^{n/2-0.25\log n}$, there are at most

$$2^{0.5\cdot 2^{n/2-.025\cdot\log n}\log n + 0.25\cdot 2^n/\sqrt{n}} \leq 2^{0.5\cdot 2^n/\sqrt{n}}$$

access structures that can be realized by a monotone span program over $\mathbb{F}_2$ of size d. On the other hand, there are more than $2^{\binom{n}{n/2}} \geq 2^{0.7\cdot 2^n/\sqrt{n}}$ access structures: We consider access structures whose minimal sets are of size exactly $n/2$ and for each set of size $n/2$, it can either be in the access structure or not in the access structure. Since

$$\lim_{n\to\infty} \frac{2^{0.5\cdot 2^n/\sqrt{n}}}{2^{0.7\cdot 2^n/\sqrt{n}}} = 0,$$

almost all access structures require shares of size greater than $0.5 \cdot 2^{n/2-0.25\log n}$ in any linear secret-sharing scheme. □

3.2.3 A Lower Bound on the Normalized Share Size in Linear Secret-Sharing Schemes

Rónyai et al. [68] proved a result similar to Claim 3.8 for every field. Note that for large q, the lower bound over $\mathbb{F}_q$ cannot be proved by counting the number of $d \times d$ matrices over $\mathbb{F}_q$. This is done by showing that over $\mathbb{F}_q$ monotone span programs of size d accept at most 2^{nd^3} access structures. The following theorem generalize the results of [9, 68].

Theorem 3.9. *For almost every access structure Γ with n parties the following property holds: For every prime-power q, the normalized share size in every linear secret-sharing scheme realizing Γ over the field $\mathbb{F}_q$ is at least $2^{n/3-o(n)}$.*

Proof. The proof is similar to the proof of Claim 3.8, with a more complex upper bound on the number of access structure that can be realized with a monotone span program of size d.

Fix some labeling function $\rho_0 : [d] \to \{p_1, \ldots, p_n\}$ and assume that there is a monotone span program $\mathcal{M} = (\mathbb{F}_q, M, \rho_0)$ accepting an access structure Γ where M is matrix over some field $\mathbb{F}_q$ of size $d \times d$. Let M_i be the i-th row of M and $M_0 = \vec{e}_1$ and define a representable matroid with a ground set $A = \{0, \ldots, d\}$ and a rank function $r(B) = \operatorname{rank}\{M_i : i \in B\}$. We next show that the rank function r together with ρ_0 determines the access structure Γ accepted by $\mathcal{M}$. Indeed, $B \in \Gamma$ if and only if $\vec{e}_1 \in \operatorname{span}\left\{M_i : p_{\rho_0(i)} \in B\right\}$ if and only if

$$\operatorname{rank}(\left\{M_i : p_{\rho_0(i)} \in B\right\}) = \operatorname{rank}(\left\{M_i : p_{\rho_0(i)} \in B\right\} \cup \{\vec{e}_1\})$$

if and only if $r(\left\{i : p_{\rho_0(i)} \in B\right\} = r(\left\{i : p_{\rho_0(i)} \in B\right\} \cup \{0\})$. Thus, the number of access structures that can be realized by a linear scheme with normalized share size is upper-bounded by the number of labeling functions ρ times the number of representable matroids with ground set $\{0, \ldots, d\}$, i.e., by $n^d \times 2^{(d+1)^3/4} \leq 2^{d^3/2}$. To conclude, for $d = 2^{n/3}/n^{1/6}$, almost all access structures do not have a linear secret-sharing scheme with normalized share size smaller than d. □

3.2.4 A Lower Bound on the Share Size in Linear Secret-Sharing Schemes with a One Bit Secret

Finally, for a one-bit secret, we obtain a lower bound of $2^{n/2-o(n)}$ on the share size of linear secret-sharing schemes over any field realizing almost all access structures even if the secret is a bit. Notice that this lower bound is on the share size (and not on the normalized share size). The constant in the exponent in Theorem 3.10 is $1/2$ (compared to a constant $1/3$ in Theorem 3.9), matching the construction of linear secret-sharing schemes for almost all access structures in Theorem 3.3 (up to lower order terms). This theorem is a special case of [4, Theorem 5.5], however, the proof of this special case is simpler.

Theorem 3.10. *For almost every access structure Γ with n parties the following property holds: For every prime-power q, the share size in every linear secret-sharing scheme over $\mathbb{F}_q$ realizing Γ with a one bit secret is at least $2^{n/2-o(n)}$.*

Proof. There are at most $n^d q^{d^2}$ monotone span programs of size d over $\mathbb{F}_q$ (as there are q^{d^2} matrices and n ways to label each row by a party). For $d > \log n$, $n^d q^{d^2} < q^{2d^2}$. The total share size in the linear secret-sharing scheme constructed from such monotone span program is $D = d \log q$. Thus, the number of linear secret-sharing schemes over $\mathbb{F}_q$ with total share size D is at most $q^{2(D/\log q)^2} < 2^{2D^2}$. Furthermore, when $q > 2^D$, the share size of each party is at least $\log q > D$ as each share contains at least on element from $\mathbb{F}_q$. Thus, the number of linear secret-sharing schemes with share size D is at most

$$\sum_{q\,:\,q\leq 2^D, q \text{ is a prime power}} 2^{2D^2} \leq 2^D \cdot 2^{2D^2} \leq 2^{3D^2}.$$

Taking $D = 0.4 \cdot 2^{n/2-0.25\log n}$, the number of access structures that have a linear secret-sharing scheme over any field with share size at most D is less than $2^{3\cdot 0.16\cdot 2^n/\sqrt{n}}$, i.e., almost all access structures require shares of size larger than D in all linear secret-sharing schemes. □

4 (G, t)-Secret-Sharing Schemes

In this section, we present a new family of schemes that we call (G, t)-secret-sharing schemes. We show that there is a close bi-directional connection between these schemes and 2-server robust CDS protocols, generalizing the connection between (non-robust) CDS protocols and forbidden graphs secret-sharing schemes. These schemes will be later used to construct graph secret-sharing schemes.

4.1 The Definition of (G, t)-Secret-Sharing Schemes

Definition 4.1. *Let $G = (V, E)$ be an undirected graph with $|V| = n$ such that $E \neq \emptyset$ and let Γ_G be the graph access structure determined by G (that is, each edge is a minimal authorized set and each independent set is forbidden). For any $0 \leq t \leq n-1$, define Γ_t as the t-out-of-n threshold access structure on V (that is, $\Gamma_t = \{A \subseteq X : |A| \geq t\}$) and define the access structure $\Gamma_{G,t}$ on V as $\Gamma_{G,t} = \Gamma_G \cup \Gamma_{t+1}$. We say a secret-sharing scheme is a* (G, t)-secret-sharing scheme *if it realizes the access structure $\Gamma_{G,t}$.*

Next, we present some properties of these schemes. If Π is a (G, t)-secret-sharing scheme, then all subsets containing edges are authorized, independent subsets of G of size at most t are forbidden, and subsets of size greater than t are authorized. If $t = 2$, then $\Gamma_{G,t}$ is a *forbidden graph* access structure determined by a graph G (for an introduction to these access structures, see [15], for example). Notice that for $t < n-1$, any (G, t)-secret-sharing scheme is also a $(G, t+1)$-secret-sharing scheme. If the size of a largest independent set of G is μ, then every subset of size $\mu + 1$ is authorized in Γ_G. Therefore, $\Gamma_{G,t} = \Gamma_G$ for every $t \geq \mu$. In particular, $\Gamma_{G,n-1} = \Gamma_G$ for every graph G.

4.2 (G, t)-Secret-Sharing Schemes from Robust CDS Protocols

We now present constructions of (G, t)-secret-sharing schemes. First, we present a transformation from robust CDS protocols to (G, t)-secret-sharing schemes. Then, using the robust CDS schemes presented in Section 2, we provide explicit (G, t)-secret-sharing schemes.

Lemma 4.2. *Let $G = (V, E)$ be a graph with $|V| = n$, and let $0 < t < n$. If there exists a 2-server t-robust CDS protocol with secrets of size m and messages of size $c(n, m)$ for functions $f : [n]^2 \to \{0, 1\}$, then there is a (G, t)-secret-sharing scheme with secrets of size m and shares of size $2 \cdot c(n, m) + \max\{m, \lceil \log n \rceil\}$.*

Proof. We construct the (G, t)-secret-sharing scheme using the scheme in Figure 2. Next we prove the correctness and privacy properties.

CORRECTNESS: Let $A \subseteq [n]$ be a minimal authorized subset in $\Gamma_{G,t}$. Then A is either in E or A is of size $t + 1$. If $A = \{i, j\}$ is in E, then $f(i, j) = 1$, i.e., the message of Alice (the first server) on i and the message of Bob (the second server) on j determines s, so the pair $\{i, j\}$ can recover s. If $|A| = t + 1$, then A can recover s using the $(t + 1)$-out-of-n secret-sharing scheme.

PRIVACY: Let A be a maximal forbidden subset. Then A does not contain any edge in E and $|A| \leq t$. The shares received from the threshold secret-sharing scheme do not provide any information about s. Now we analyze the information provided by the messages of $\mathcal{P}$. The parties in A receive Alice's messages for A and Bob's messages for A. Observe that the set $A \times A$ does not contain edges of G, thus, $A \times A$ is a zero-set of f and the t-robustness of $\mathcal{P}$ guarantees the privacy of the scheme.

The maximum share size of the resulting scheme is twice the message size of $\mathcal{P}$ plus the share size of the $(t + 1)$-out-of-n secret-sharing scheme. □

The secret: An element $s \in S$.
The parties: $V = \{1, \ldots, n\}$.
The access structure: $\Gamma_{G,t}$ for some graph $G = (V, E)$ and $0 \leq t \leq n - 1$.
The scheme:

- Let $f : [n] \times [n] \to \{0, 1\}$ be the function defined as $f(i, j) = 1$ if and only if $(i, j) \in E$.
- Let $\mathcal{P}$ be a 2-server t-robust CDS protocol with secrets from $\{0, 1\}^m$ for the function f; denote its servers by Alice and Bob.

Then,

1. Execute the protocol $\mathcal{P}$ for the secret s.
2. Share s independently among V with a $(t + 1)$-out-of-n secret-sharing scheme.
3. The share of party $i \in V$ is the message of Alice on the input i, the message of Bob on the input i, and the share of i in the $(t + 1)$-out-of-n secret-sharing scheme.

Figure 2: A (G, t)-secret-sharing scheme Π for a graph $G = (V, E)$.

In Lemma 4.2, we showed a way to construct (G, t)-secret-sharing schemes from t-robust CDS protocols. Conversely, we can also construct robust CDS protocols from (G, t)-secret-sharing schemes, as shown in Lemma 4.3.

Lemma 4.3. *Let $f : [n] \times [n] \to \{0,1\}$ be a function and let $0 < t < n$. Define $G = (([n] \times \{1\}) \cup ([n] \times \{2\}), E)$ as the bipartite graph with $E = \{((i,1),(j,2)) : i \in [n], j \in [n], f(i,j) = 1\}$. If there exists a $(G, 2t)$-secret-sharing scheme with secrets of size m and maximum share size $c(2n, m)$, then there exists a 2-server t-robust CDS protocol for f with message size $c(2n, m)$.*

Proof. Let Π be a $(G, 2t)$-secret-sharing scheme. We define a 2-server t-robust CDS protocol $\mathcal{P}$ for f as follows. The message spaces M_1 and M_2 of the servers are the spaces of shares of parties $[n] \times \{1\}$ and $[n] \times \{2\}$, respectively. The common randomness r is the randomness of the dealer in Π. The function $\textsc{Enc}_i(j, s, r)$ for $i \in \{1, 2\}$ outputs the share of party (j, i) with the secret s and randomness r, and $\textsc{Dec}$ is the reconstruction function of Π.

The correctness of $\mathcal{P}$ is guaranteed because every pair in E is authorized in Π. The t-robustness of $\mathcal{P}$ is guaranteed because every zero-set $Z_1 \times Z_2$ where $|Z_1|, |Z_2| \leq t$ corresponds to an independent set $(Z_1 \times \{1\}) \cup (Z_2 \times \{2\})$ of size at most $2t$ in G, thus the messages of the inputs in $Z_1 \cup Z_2$ are shares of a forbidden set in Π. □

Now that we showed the connection between (G, t)-secret-sharing schemes from t-robust CDS protocols, we present (G, t)-secret-sharing schemes that use Theorem 2.12 and Theorem 2.11.

Lemma 4.4. *Let $G = (V, E)$ be a graph with $|V| = n$, and let $1 \leq t < n/2$. If there exist a 2-server CDS protocol with message size $c(n, m)$ for functions with domain size n and secrets of size m, then there exists a (G, t)-secret-sharing scheme with maximum share size $O(t^2 \log^3 t \log^2 n \cdot c(n, m))$, and a (G, t)-secret-sharing scheme with secrets of size $\Theta(mt \log t \log^2 n)$ and normalized maximum share size $O(t \log^2 t \cdot c(n, m)/m)$.*

Proof. Theorem 2.12 guarantees that there exists a 2-server t-robust CDS protocol with message size $\ell(n) = O(t^2 c(n, m) \log^3 t \log^2 n)$, and a 2-server t-robust CDS protocol with secrets of size $m' = \Theta(mt \log t \log^2 n)$ with normalized message size $\ell(n)/m' = O(t \log^2 t \cdot c(n, m)/m)$. Using these 2-server t-robust CDS protocols and Lemma 4.2 we obtain the lemma. □

We conclude this section presenting different (G, t)-secret-sharing schemes that are obtained from robust CDS schemes applying Lemma 4.2 and Lemma 4.4.

Theorem 4.5. *Let $G = (V, E)$ be a graph with $|V| = n$ and let $1 < t < n$.*

1. *There exists a (G, t)-secret-sharing scheme with moderately-short secrets of size $O(t \log^3 n)$, normalized maximum share size*

$$n^{O(\sqrt{\log \log n / \log n})} t \log^2 n = n^{o(1)} t \log^2 n,$$

   and normalized total share size $n^{1+O(\sqrt{\log \log n / \log n})} t \log^2 n = n^{1+o(1)} t \log^2 n$;

2. *For every prime power q, there exists a linear (G, t)-secret-sharing scheme over $\mathbb{F}_q$ with and maximum share size*

$$O\left((t \log^2 t + \sqrt{n}) t \log t \log^2 n \log q\right);$$

3. *There exists an integer p_0 such that for every prime power $q > p_0$, there exists a multi-linear (G, t)-secret-sharing scheme over $\mathbb{F}_q$ with moderately-short secrets of size $\Theta(t^2 \log t \log^2 n \log n \log q)$ and normalized maximum share size $O(t \log^2 t + \sqrt{n})$;*

4. *There exists a multi-linear (G,t)-secret-sharing scheme over $\mathbb{F}_2$ with secrets of size 2^{n^2} and normalized maximum share size $O(t\log^2 t)$.*

Proof. Scheme 1: By Theorem 2.7, for any function $f : [n]^2 \to \{0,1\}$ there exists a 2-server CDS protocol with secret of size $m = 1$ and messages size $c(n,1) = n^{O(\sqrt{\log\log n/\log n})}$. Applying Theorem 2.12 with the CDS protocol from Theorem 2.7 results in a 2-server t-robust CDS protocol with secrets of size $O(t\log t\log^2 n) = O(t\log^3 n)$, message size $O(n^{O(\sqrt{\log\log n/\log n})}t^2\log^5 t)$, and normalized message size $O(n^{O(\sqrt{\log\log n/\log n})}t\log^2 t)$. By Lemma 4.2, there is a (G,t)-secret-sharing with secrets of size $O(t\log^3 n)$ and maximum share size $O(n^{O(\sqrt{\log\log n/\log n})}t^2\log^5 t)$, thus with normalized maximum share size $O(n^{O(\sqrt{\log\log n/\log n})}t\log^2 n)$ and with normalized total share size $O(n^{1+O(\sqrt{\log\log n/\log n})}t\log^2 n)$.

Scheme 2: Theorem 2.11 guarantees that for $t \le n/(2\log^2 n)$ there exists a linear 2-server t-robust CDS protocol over $\mathbb{F}_q$ with message size $O\left((t\log^2 t+\sqrt{n})t\log t\log^2 n\log q\right)$. Thus, by Lemma 4.2 there is a (G,t)-secret-sharing scheme where the maximum share size is the above message size. For $t > n/(2\log^2 n)$, the upper bound also holds because there is always a linear (G,t)-secret-sharing with maximum share size $O(n/\log n)$.

Scheme 3: Theorem 2.11 also guarantees, for a large enough q, a 2-server (t,n)-robust CDS protocol with secrets of size $\Theta(t^2\log t\log^2 n\log q)$ and normalized message size $O(t\log^2 t+\sqrt{n})$. Again, we construct the desired (G,t)-secret-sharing with from the robust CDS protocol applying Lemma 4.2.

Scheme 4: By Theorem 2.6, there exists a multi-linear CDS protocol over $\mathbb{F}_2$ with normalized message size $c(n,m)/m = 3$ for secrets of size 2^{n^2}. Applying Lemma 4.4, we obtain a multi-linear (G,t)-secret-sharing over $\mathbb{F}_2$ with normalized maximum share size $O(t\log^2 t\cdot c(n,m)/m) = O(t\log^2 t)$. □

5 Secret-sharing Schemes for Almost All Graphs

In this section we study the maximum share size of secret-sharing schemes for almost all graphs and for almost all graphs in $\mathscr{G}(n,p)$ for different values of p. The previous and new results for almost all graphs are summarized in Figure 1, while the results for $\mathscr{G}(n,p)$ are summarized in Figure 5.

Schemes presented in this section rely on the properties of almost all graphs shown in Section 2.4, and use the (G,t)-secret-sharing schemes presented in Section 4. In order to understand the share size of secret-sharing schemes for almost all graphs, we provide lower bounds for them in Theorem 5.5 and Theorem 5.7.

5.1 Schemes for Almost all Graphs

As a consequence of Lemma 2.13, the size of every independent set in almost every graph in $\mathscr{G}_n$ is $O(\log n)$. We observed in Section 4 that a (G,t)-secret-sharing scheme is also a secret-sharing scheme realizing G when t is bigger than the size of a largest independent set of G. Hence, we consider the four constructions presented in Theorem 4.5 for $t = O(\log n)$. In Theorem 5.1 we present the resulting schemes.

Theorem 5.1. *Almost all graphs with n vertices can be realized by the following schemes.*

1. *A secret-sharing scheme with maximum share size $n^{O(\sqrt{\log\log n/\log n})} = n^{o(1)}$,*
2. *A linear secret-sharing scheme over $\mathbb{F}_q$ with maximum share size $\tilde{O}(\sqrt{n}\log q)$ for every prime power q,*
3. *A multi-linear secret-sharing scheme over $\mathbb{F}_q$ with normalized maximum share size $O(\sqrt{n})$ and moderately-short secrets of size $\Theta(\log q\log^3 n\log\log n)$ for a large enough q, and*

4. *A multi-linear secret-sharing scheme over $\mathbb{F}_2$ with normalized maximum share size $O\left(\log n(\log\log n)^2\right)$ for secrets of size 2^{n^2}.*

5.2 Secret-sharing Schemes for $\mathscr{G}(n,p)$

In order to study properties of sparse graphs, we study $\mathscr{G}(n,n^{-\alpha})$ for a constant $0<\alpha<1$. Almost all graphs in $G(n,n^{-\alpha})$ have maximal independent sets of size at most $t=O(n^{\alpha}\log n)$. Following the procedure we developed in the previous section, we can construct secret-sharing schemes for almost all graphs in $G(n,n^{-\alpha})$ using Theorem 4.5. Similar bounds can be obtained for linear schemes and multi-linear schemes. They are presented in Figure 5.

Theorem 5.2. *Let $0<\alpha<1$ be a constant. Almost every graph in $\mathscr{G}(n,n^{-\alpha})$ can be realized by a secret-sharing scheme with normalized maximum share size $n^{\min(\alpha,1-\alpha)+o(1)}$ and secret of size $\tilde{O}(\sqrt{n})$.*

Proof. We present two schemes Π_1 and Π_2 for almost all graphs in $\mathscr{G}(n,n^{-\alpha})$. The scheme Π_1 consists on sharing the secret for each edge independently. By Lemma 2.15, almost every graph in $\mathscr{G}(n,n^{-\alpha})$ has maximum degree of at most $2n^{1-\alpha}$. Therefore, the maximum share size of Π_1 is $2n^{1-\alpha}$ for almost all graphs in $\mathscr{G}(n,n^{-\alpha})$.

The second scheme Π_2 is obtained from Theorem 4.5. For almost every graph in $\mathscr{G}(n,n^{-\alpha})$ the size of a maximum independent set is $O(n^{\alpha}\log n)$ (by Lemma 2.14). Thus, we let Π_2 be the $(G,O(n^{\alpha}\log n))$-secret-sharing scheme of Theorem 4.5 with secret of size $\Theta(t\log^3 n)=\Theta(n^{\alpha}\log^4 n)$ and normalized maximum share size $O(n^{o(1)}t\log^2 n)=O(n^{\alpha+o(1)}\log^3 n)=n^{\alpha+o(1)}$.

Therefore, almost every graph in $\mathscr{G}(n,n^{-\alpha})$ can be realized by a secret-sharing scheme with normalized maximum share size $\min(2n^{1-\alpha},n^{\alpha+o(1)})\leq n^{\min(1-\alpha,\alpha)+o(1)}$. □

For $\alpha\leq 1/2$, the best choice is Π_1, and for $\alpha>1/2$, the best choice is Π_2. For $\alpha=1/2$, the normalized maximum share size of almost all graphs in $\mathscr{G}(n,n^{-\alpha})$ in our scheme is $O(\sqrt{n})$. This is the constant α that gives the worst upper bound on the normalized maximum share size of secret-sharing schemes for $\mathscr{G}(n,n^{-\alpha})$.

Finally, we study properties of very dense graphs by analyzing $\mathscr{G}(n,1-n^{-\alpha})$ for a constant $0<\alpha<1$. By Lemma 2.14, the size of a maximum independent set for almost all graphs in $\mathscr{G}(n,1-n^{-\alpha})$ is constant. As we saw above, graphs with small independent sets admit more efficient schemes. In Theorem 5.4 we present secret-sharing schemes for all graphs in $\mathscr{G}(n,1-n^{-\alpha})$. Two of the schemes we present in Theorem 5.4 follow quite easily from our previous results. In contrast, the linear scheme we construct in Theorem 5.4 does not follow from previous results on robust CDS protocols. Rather, it follows from the following theorem of [15] on the total share size for forbidden graph secret sharing schemes and the techniques of [5].

Theorem 5.3 ([15, Theorem 6])**.** *Let $G=(V,E)$ graph with n vertices and at least $\binom{n}{2}-n^{1+\beta}$ edges, for some $0\leq\beta<1$. Then for every prime-power $q>n$ there is a linear $(G,2)$-secret-sharing scheme over $\mathbb{F}_q$ that with total share size $\tilde{O}(n^{1+\beta/2}\log q)$*

Theorem 5.4. *Let $0\leq\beta<1$ be a constant. Almost all graphs in $\mathscr{G}(n,1-n^{\beta-1})$ can be realized by a secret-sharing scheme with maximum share size $n^{o(1)}$, a linear secret-sharing scheme over $\mathbb{F}_q$ with total share size $\tilde{O}(n^{1+\beta/2}\log q)$ for every prime-power $q>n$, and a multi-linear secret-sharing scheme over $\mathbb{F}_2$ with exponentially long secrets of size 2^{n^2} and normalized maximum share size $O(1)$.*

Proof. By Lemma 2.14, the size of a maximum independent set for almost all graphs in $\mathscr{G}(n, 1-n^{-\alpha})$ is some constant c. The non-linear secret-sharing scheme and the secret-sharing scheme with long secrets are obtained by applying Theorem 4.5 with $t = O(1)$.

To construct the linear secret-sharing scheme we note that the maximum degree of almost every graph G in $\mathscr{G}(n, 1-n^{\beta-1})$ is at least $n-2n^{\beta}$ (by Lemma 2.15 applied to $\overline{G}$), thus the number of edges in G is at least $\binom{n}{2}-n^{1+\beta}$. The linear scheme is derived by using the technique of [5] to transform the $(G,2)$-secret-sharing scheme from Theorem 5.3 to a (G,c)-secret-sharing scheme: Let $\mathcal{H} = \left\{h_i : [n] \to [c^2] : 1 \leq i \leq \ell\right\}$ be a family of perfect hash functions,[3] where $|\mathcal{H}| = \ell = O(\log n)$. The (G,c)-secret-sharing scheme, denoted Π, is as follows:

- Input: a secret $s \in \mathbb{F}_q$.
- Choose $\ell - 1$ random elements $s_1, \ldots, s_{\ell-1}$ from $\mathbb{F}_q$ and let $s_\ell = s - (s_1 + \cdots + s_{\ell-1})$.
- For every $i \in \{1, \ldots, \ell\}$ and every $a, b \in \left\{1, \ldots, c^2\right\}$, independently share s_i using the $(G,2)$-secret-sharing scheme and give the share of vertex v to v if and only if $h_i(v) \in \{a, b\}$.

For the correctness of the scheme Π, let (u,v) be an edge in G (i.e., an authorized set). For every i, the parties u, v can reconstruct s_i from the scheme for $a = h(u), b = h(v)$. For the privacy of Π, let B be an independent set in G (i.e., a forbidden set). By Lemma 2.14, we can assume that the size of B is at most c, thus, there exists a hash function $h_i \in \mathcal{H}$ such that $h_i(u) \neq h_i(v)$ for every distinct $u, v \in B$. Therefore, in any sharing of s_i for some values a, b the parties in B hold at most 2 shares, and these shares are of a forbidden set. The privacy of the $(G,2)$-secret-sharing scheme implies that the parties in B do not get any information on s_i from this execution. Since all executions of the $(G,2)$-secret-sharing scheme are executed with an independent random string, the parties in B do not get any information on s_i from the shares of Π, hence they get no information on s. Note that the total share size in Π is $O(\log n)$ times the total share size of the $(G,2)$-secret-sharing scheme. □

5.3 Lower Bounds for the Share Size for Almost All Graphs

Next, we present lower bounds for the maximum share size of secret-sharing schemes for almost all graphs. This question was first addressed by Csirmaz in [39], where he proved a lower bound which we include in Theorem 5.5.

Theorem 5.5. *For almost every graph G, the normalized maximum share size of every secret-sharing scheme realizing G is $\Omega(\log\log n)$, and the normalized maximum share size of every multi-linear secret-sharing scheme realizing G is $\Omega(\log^{1/2} n)$.*

Proof. Both bounds are a consequence of Lemma 2.16 (which says that almost all n-vertex graphs contain all graphs of size $\log n$ as an induced graph), taking different graphs with $\log n$ vertices. The first bound was proved by Csirmaz in [39], taking the family of hypercube graphs (or the graphs of [45]). The ℓ-cube has 2^ℓ vertices, and its normalized maximum share size is at least $\Omega(\ell)$ [39]. Hence, taking $\ell = \lfloor \log\log n \rfloor$, almost every graph with n vertices contains as an induced graph the ℓ-cube, which requires normalized maximum share size $\Omega(\log\log n)$.

[3] A family $\mathcal{H}$ is a family of perfect hash functions for sets of size at most c if for every $B \subset \{1, \ldots, n\}$ such that $|B| \leq c$, there exists a function $h \in \mathcal{H}$ such that h is one-to-one on B, that is, $h(u) \neq h(v)$ for every distinct $u, v \in B$. By a standard probabilistic argument, such family of size $O(c \log n)$ exists. For a constant c, the size of the family is $O(\log n)$.

The second bound is a consequence of the results in [16, 11]. By [16], for every r there exists a graph H_r where H_r has r vertices and for every finite field $\mathbb{F}_q$ the normalized maximum share size in any linear secret-sharing scheme realizing H_r over $\mathbb{F}_q$ is $\Omega(r^{1/2})$. By [11], this bound applies to multi-linear secret-sharing schemes as well. Hence, the normalized maximum share size of every multi-linear secret-sharing scheme realizing H_r is also $\Omega(r^{1/2})$. By Lemma 2.16, almost every graph G with $n = r^2 2^{r/2}$ vertices contains H_r as an induced graph, hence the normalized maximum share size in any multi-linear secret-sharing scheme over $\mathbb{F}_q$ realizing G is $\Omega(r^{1/2}) = \Omega(\log^{1/2} n)$. □

Remark 5.6. Lemma 2.16 provides a connection between the maximum share size of schemes for *every* graph access structure with $r = \log n$ vertices and the maximum share size of schemes for almost all graph access structures with n vertices. In Theorem 5.5 we used it in one direction, but it could also be used in the converse direction. For instance: if there exist secret-sharing schemes for almost all n-vertex graphs with (normalized) maximum share size $\ll \frac{\log n}{\log \log n}$, then there exist secret-sharing schemes realizing every r-vertex graph with (normalized) maximum share size $\ll r/\log r$, which is currently the best upper bound [47].

In Theorem 5.7, we quote a lower bound on the maximum share size for linear graph secret-sharing schemes, proved in [63, 14]. Notice, however, that this bound does not grow as a function of the size of the secrets.

Theorem 5.7 ([63, 14])**.** *For almost every graph G, the maximum share size of every linear secret-sharing scheme realizing G is $\Omega(\sqrt{n})$.*

6 Secret-sharing Schemes for Very Dense Graphs

In this section we study secret-sharing schemes for very dense graphs, i.e., graphs with n vertices and at least $\binom{n}{2} - n^{1+\beta}$ edges for some $0 \leq \beta < 1$. This problem was originally studied in [13], and the best previously known upper bounds on the maximum share size and the total share size are presented in Theorems 6.1 and 6.2. The normalized total share size of the schemes presented in this section is smaller than the maximum share size of previous schemes.

Theorem 6.1 ([13])**.** *Let $G = (V, E)$ be a graph with $|V| = n$ and $|E| \geq \binom{n}{2} - n^{1+\beta}$ for some $0 \leq \beta < 1$. Then, there exists a linear secret-sharing scheme realizing G with maximum share size $\tilde{O}(n^{1/2+\beta/2})$, total share size $\tilde{O}(n^{3/2+\beta/2})$, and secret of size $O(\log n)$.*

The above theorem hides poly-logarithmic factors in the share size. It was also shown in [13] that these poly-logarithmic factors can be avoided if we consider multi-linear secret-sharing schemes and normalized share size: for the graphs considered in Theorem 6.1, there exists a multi-linear secret-sharing scheme with normalized maximum share size $O(n^{1/2+\beta/2})$ and secret of size $O(\log^2 n)$.

In [13], there is another secret-sharing construction for very dense graphs, presented in Theorem 6.2. The total share size of this scheme is smaller than the one in Theorem 6.1, but the maximum share size may be larger.

Theorem 6.2 ([13])**.** *Let $G = (V, E)$ be a graph with $|V| = n$ and $|E| \geq \binom{n}{2} - n^{1+\beta}$ for some $0 \leq \beta < 1$. There exists a linear secret-sharing scheme realizing G with total share size $\tilde{O}(n^{5/4+3\beta/4})$.*

As an observation, notice that as a direct implication of the results in previous sections we can construct a scheme whose maximum share size is similar to the maximum share size as in the scheme of Theorem 6.2. The size of the largest independent set in a graph G with at least $\binom{n}{2} - n^{1+\beta}$ edges is at most $t = \sqrt{n^{1+\beta}}$. Then, by Theorem 4.5, there exists a multi-linear (G, t)-secret-sharing scheme with secrets of size $O(n^{1/2+\beta/2} \log^4 n)$ and normalized maximum share size

$$O(t \log^2 t + \sqrt{n}) = O(n^{1/2+\beta/2} \log^2 n).$$

Notice, however, that this result is weaker than the one in Theorem 6.2, because in this new scheme the secret is larger, and the normalized maximum share size is bigger by a factor of $\log^2 n$.

6.1 New Constructions

We use (G, t)-secret-sharing schemes, described in the Section 4, to construct secret-sharing schemes for *all* very dense graphs. Our main result for dense graphs is Theorem 6.4, where we show that graphs with at least $\binom{n}{2} - n^{1+\beta}$ edges admit secret-sharing schemes with normalized total share size $n^{1+\beta+o(1)}$. This result nearly matches the best total share size for sparse graphs with at most $n^{1+\beta}$ edges (for which we share the secret independently for each edge). The construction follows the ideas described in the introduction.

The secret: An element $s \in S$.
The parties: $V = \{1, \ldots, n\}$
The scheme:

1. Let $\beta < \alpha < (1 + \beta)/2$ and $n' = n^{1+\beta-\alpha}$.
2. Let $A \subseteq V$ be a subset of n' vertices of lowest degree and $G' = (A, E \cap (A \times A))$.
3. Share s among A using Π_1, a secret-sharing scheme realizing G'.
4. Choose $r \in S$ uniformly at random.
5. Share r using Π_2, a $(G, 2n^\alpha + 1)$-secret-sharing scheme.
6. Share $r + s$ using Π_3, a secret-sharing scheme where A is the only maximal forbidden subset (that is, give s to every party not in A).

Figure 3: A secret-sharing scheme Π_{dense} realizing a graph $G = (V, E)$ with $|E| \geq \binom{n}{2} - n^{1+\beta}$ for some $0 \leq \beta < 1$.

Lemma 6.3. *Let $G = (V, E)$ be a graph with $|V| = n$ and $|E| \geq \binom{n}{2} - n^{1+\beta}$ for some $0 \leq \beta < 1$. The scheme described in Figure 3 is a secret-sharing scheme realizing G.*

Proof. Let A be a subset n' vertices of lowest degree in G, let v be a vertex of highest degree in A, and denote its degree by $n - d - 1$, i.e., each vertex in A misses at least d edges (and each missed edge touches at most two vertices in A). Then $2n^{1+\beta} \geq d \cdot |A| = d \cdot n^{1+\beta-\alpha}$, so $d \leq 2n^\alpha$. In particular, every vertex in $V \setminus A$ misses at most d edges, and so the size of every independent set containing at least one vertex in $V \setminus A$ is at most than $d + 1 \leq 2n^\alpha + 1$.

Next we prove the correctness and privacy of the scheme in Figure 3.

CORRECTNESS: The minimal authorized subsets in Γ_G are the edges of the graph. Let (x, y) be an edge of G. If $\{x, y\} \subseteq A$, then $\{x, y\}$ is authorized in Π_1. Else, it is authorized in Π_2 and in Π_3, and so $\{x, y\}$ can recover the secret.

PRIVACY: Let B be a forbidden set of Γ_G, that is, B is an independent set of G. In particular, B is an independent set of G' and so it does not learn any information on s from Π_1. If $B \subseteq A$, then B is forbidden in Π_3 because it is contained in A, hence it does not learn any information on $s + r$ from Π_3 and does not learn any information on s from Π_2, regardless of the information it may learn on r. If B is not contained in A, then it is an independent set of G of size at most $2n^\alpha + 1$. Hence B does not learn any information on r in Π_2 and does not learn any information on s from Π_3, regardless of the information it may learn on $r + s$. □

In Theorem 6.4, we use Π_{dense} recursively to obtain our improved secret-sharing scheme for dense graphs. In order to improve the readability of the proof of Theorem 6.4, we say that a scheme has share size exponent γ when the normalized total share size is $n^{\gamma+o(1)}$.

Theorem 6.4. *Let $G = (V, E)$ be a graph with $|V| = n$ and $|E| \geq \binom{n}{2} - n^{1+\beta}$ for some $0 \leq \beta < 1$. Then G can be realized by a secret-sharing schemes with secrets of size $O(n \log^3 n)$ and normalized total share size $n^{1+\beta+o(1)}$.*

Proof. We show that there exists a sequence of families of schemes $\{\mathcal{F}_k\}_{1 \leq k \leq \log n}$ that for any β and n realize graphs with n vertices and at least $\binom{n}{2} - n^{1+\beta}$ edges such that the scheme $\mathcal{F}_k$ uses the scheme $\mathcal{F}_{k-1}$ to reduce the share size and the total share size of $\mathcal{F}_{\log n}$ is $n^{1+\beta+o(1)}$. Formally, we show that for $k \geq 1$, for every n and $0 < \beta < 1$ there exist a secret-sharing scheme $\mathcal{F}_k$ with share size exponent $\frac{2k+1}{2k} + \frac{2k-1}{2k}\beta$ for every graph with at least $\binom{n}{2} - n^{1+\beta}$ edges.

The proof is by induction. Theorem 6.1 shows that it holds for $k = 1$, that is, there is a scheme with share size exponent $3/2 + \beta/2$. Now suppose that for every n and β there exist a secret-sharing scheme $\mathcal{F}_k$ with share size exponent $\frac{2k+1}{2k} + \frac{2k-1}{2k}\beta$ for every graph with at least $\binom{n}{2} - n^{1+\beta}$ edges. We will show that for every n and β there exists a secret-sharing scheme $\mathcal{F}_{k+1}$ with share size exponent $\frac{2(k+1)+1}{2(k+1)} + \frac{2(k+1)-1}{2(k+1)}\beta$ for every graph with at least $\binom{n}{2} - n^{1+\beta}$ edges.

We construct the secret-sharing scheme $\mathcal{F}_{k+1}$ using the secret-sharing scheme described in Figure 3. We combine a scheme Π_1 for very dense graphs, a scheme Π_2 for graphs with a bound on the size of their maximal independent set, and a scheme Π_3. As described in Figure 3, the scheme Π_3 has maximum share size 1 and total share size $n - |A|$ (that is, it gives $r \oplus s$ to every party not in A).

Next, we present the schemes Π_1 and Π_2. Let $\beta' = \frac{\alpha}{1+\beta-\alpha}$. The graph G' has at least $\binom{n'}{2} - n^{1+\beta} = \binom{n'}{2} - \left((n')^{1/(1+\beta-\alpha)}\right)^{1+\beta} = \binom{n'}{2} - (n')^{1+\beta'}$ edges. By the induction hypothesis, we can take the secret-sharing scheme $\mathcal{F}_k$ as Π_1 (i.e., as the secret-sharing scheme realizing G'); the total share size of $\mathcal{F}_k$ is $\tilde{O}((n')^\delta)$ with $\delta = 1 + \frac{1}{2k} + \frac{2k-1}{2k}\beta'$. Since $n' = n^{(1+\beta-\alpha)}$, this scheme has share size exponent

$$(1 + \beta - \alpha)\left(\frac{2k+1}{2k} + \frac{2k-1}{2k}\frac{\alpha}{1+\beta-\alpha}\right) = \frac{2k+1}{2k} + \frac{2k+1}{2k}\beta - \frac{2\alpha}{2k}.$$

We use the first scheme of Theorem 4.5 to realize the $(G, 2n^\alpha + 1)$ secret-sharing scheme Π_2 with secrets of size $O(n^\alpha \log^3 n) \leq O(n \log^3 n)$ and share size exponent $1 + \alpha$.

The resulting construction has share size exponent

$$\max\left\{1 + \alpha, \frac{2k+1}{2k} + \frac{2k+1}{2k}\beta - \frac{2\alpha}{2k}\right\}.$$

In order to balance these two terms, we take

$$\alpha = \frac{1}{2k+2}(1+(2k+1)\beta) = \frac{1}{2k+2} + \frac{2k+1}{2k+2}\beta.$$

Then the resulting scheme has share size exponent

$$1+\alpha = \frac{2k+3}{2k+2} + \frac{2k+1}{2k+2}\beta,$$

which concludes the proof of the induction.

We claim that if we take the above sequence $\{\mathcal{F}_k\}_{1\leq k\leq \log(n)/2}$, then $\mathcal{F}_{\log(n)/2}$ has share size exponent $1+\beta$. Notice that for $\ell = \log(n)/2$

$$n^{\frac{2\ell+1}{2\ell}+\frac{2\ell-1}{2\ell}\beta} = n^{1+\beta+O(1/\log n)} = O(n^{1+\beta}).$$

The secret: An element $s \in S$.
The parties: $V = \{1,\ldots,n\}$
The scheme:

1. Let $\ell = \frac{\log n}{2}$, $V_\ell = V$, $n_\ell = n$, and $\beta_\ell = \beta$.
2. For $k = \ell$ downto 1 do:
   (a) Let $G_k = (V_k, E \cap V_k \times V_k)$
   ($*$ G_k is a graph with n_k vertices and at least $\binom{n_k}{2} - n_k^{1+\beta_k}$ edges $*$)
   (b) Let $\alpha_k = \frac{1}{2k+2} + \frac{2k+1}{2k+2}\beta_k$, $n_{k-1} = n_k^{1+\beta_k-\alpha_k}$, and $\beta_{k-1} = \frac{\alpha_k}{1+\beta_k-\alpha_k}$.
   (c) Let $V_{k-1} \subseteq V_k$ be the subset of size n_{k-1} with the vertices of lowest degree in G_k.
   (d) Choose uniformly at random an element $r_k \in S$.
   (e) Share r_k with a $(G_k, 2n_k^{\alpha_k}+1)$-secret-sharing scheme.
   (f) Give $r_k + s$ to every party in $V_k \setminus V_{k-1}$.
3. Share the secret s with the secret-sharing scheme of Theorem 6.1 realizing the graph $G_0 = (V_0, E \cap V_0 \times V_0)$.

Figure 4: A secret-sharing scheme $\mathcal{F}_{\frac{\log n}{2}}$ for a graph $G = (V,E)$ with $|E| \geq \binom{n}{2} - n^{1+\beta}$ for some $0 < \beta < 1$.

For the clarity of the presentation (and to ensure that we do not misuse the asymptotic analysis), we explicitly describe $\mathcal{F}_{\log(n)/2}$ in Figure 4. In $\mathcal{F}_{\log(n)/2}$ we execute a $(G_k, 2n_k^{\alpha_k}+1)$-secret-sharing scheme for $k = \frac{\log n}{2}$ downto 1. The normalized total share size in the $(G_k, 2n_k^{\alpha_k}+1)$-secret-sharing scheme is $n_k^{1+\alpha_k+o(1)}$. It can be check that $n_{k-1}^{1+\alpha_{k-1}} = n_k^{1+\alpha_k}$ for every k, thus, the total share size is

$$O\left(\log n \cdot n_\ell^{1+\alpha_\ell+o(1)}\right) = O\left(\log n \cdot n^{1+\beta+o(1)}\right).$$

□

	total inf. ratio (mod. short secrets)	total share size of linear schemes over $\mathbb{F}_q$	total inf. ratio multi-linear schemes
$n^{1+\beta}$ edges	$n^{1+\beta}$ [55] $\Omega(n \log n)$ [40]	$n^{1+\beta} \log q$ [55] $\Omega(n^{\min(1+\beta,3/2)} \log q)$ [16]	$n^{1+\beta}$ [55] $\Omega(n^{\min(1+\beta,3/2)})$ [16, 11]
$\mathscr{G}(n, n^{\beta-1})$	$n^{\min(1+\beta,2-\beta)+o(1)}$ Th. 5.2	$n^{\min(1+\beta,3-2\beta)+o(1)} \log q$ Th. 5.2	$n^{\min(1+\beta,2-\beta)}$ Th. 5.2
$\binom{n}{2} - n^{1+\beta}$ edges	$n^{1+\beta+o(1)}$ Th. 6.4 $\Omega(n \log n)$ [13]	$\tilde{O}(n^{5/4+\beta/4} \log q)$ [13] $\Omega(n^{1+\beta/2} \log q)$ [13]	$\tilde{O}(n^{1+\beta})$ Rem. 6.5 $\Omega(n^{1+\beta/2})$ [13, 11]
$\mathscr{G}(n, 1 - n^{\beta-1})$	$n^{1+o(1)}$ Th. 5.4	$O(n^{1+\beta/2}) \log q$ Th. 5.4	$O(1)$ Th. 5.4

Figure 5: Total share size for different families of graphs and constant $0 < \beta < 1$. Note that almost all graphs in $\mathscr{G}(n, n^{\beta-1})$ and in $\mathscr{G}(n, 1 - n^{\beta-1})$ have $\Theta(n^{1+\beta})$ and $\binom{n}{2} - \Theta(n^{1+\beta})$ edges, respectively

Remark 6.5. In Theorem 6.4, we combine the secret-sharing scheme for very dense graphs in Theorem 6.1 with several instances of the first scheme of Theorem 4.5. Instead, if we replace the former by the fourth scheme of Theorem 4.5, we obtain a multi-linear secret-sharing scheme with secrets of exponential size and normalized total share size $\tilde{O}(n^{1+\beta})$ for exponentially long secrets.

We can also replace it by other schemes presented described in this paper; the resulting schemes improve the current best schemes just for some values of β. For $\beta > 2/5$, using the third scheme of Theorem 4.5 in just one recursion step, we get a multi-linear scheme with moderately-short secrets of size $\tilde{O}(n^2)$ and total share size $n^{7/6+5\beta/6+o(1)}$. For $\beta < 1/3$, combining Theorem 5.3 and Lemma 4.4 and one recursion step, we get a scheme with normalized total share size $\tilde{O}(n^{7/6+\beta})$. These optimizations for specific values of β are not presented in Figure 5.

Remark 6.6. In Figure 5, we present lower bounds on the total share size for graphs with at most $n^{1+\beta}$ edges and graphs with at least $\binom{n}{2} - n^{1+\beta}$ edges. The lower bounds for very dense graphs were presented in [13]. For sparse graphs the lower bound follows from similar arguments as we next explain. We prove that for every $0 \leq \beta \leq 1$ and every n there exists a graph with n vertices and $n^{1+\beta}$ edges such that the normalized total share size of any linear and multi-linear secret-sharing scheme realizing the graph is $\Omega(n^{1+\beta})$. The starting point is result of [16], constructing for every n a graph H_n with n vertices and $n^{3/2}$ edges such that the normalized total share size of linear secret-sharing schemes realizing the graph is $\Omega(n^{3/2})$. By [11], the same lower bound holds for multi-linear schemes. This proves the claim for $\beta = 1/2$.

For $0 < \beta < 1/2$, we construct a graph $G_{n,\beta}$ with n vertices and $n^{1+\beta}$ edges that requires normalized total share size $\Omega(n^{1+\beta})$ for linear and multi-linear schemes. We partition the set of vertices into $n^{1-2\beta}$ parts of size $n^{2\beta}$, and in each part we construct a copy of the graph $H_{n^{2\beta}}$ of [16]; there are no edges between the parts. In the graph $H_{n^{2\beta}}$ there are $n^{3\beta}$ edges and it requires normalized total share size $\Omega(n^{3\beta})$. Since the graph $G_{n,\beta}$ contains $n^{1-2\beta}$ copies of $H_{n^{2\beta}}$, the graph $G_{n,\beta}$ contains $n^{1-2\beta} n^{3\beta} = n^{1+\beta}$ edges and require normalized total share size $\Omega(n^{1+\beta})$.

For $1/2 < \beta < 1$, we construct a graph $G_{n,\beta}$ with n vertices and $n^{1+\beta}$ edges that requires normalized total share size $\Omega(n^{3/2})$ for linear and multi-linear schemes. Informally, we need to add edges to H_n while maintaining the lower bound. To achieve this goal, we partition the vertices of $G_{n,\beta}$ to two equal parts of size $n/2$. In one part we construct a copy of H_n. This part contains $(n/2)^{3/2}$ edges and its normalized total share size is at least $\Omega(n^{3/2})$, hence normalized total share size of $G_{n,\beta}$ is also at least $\Omega(n^{3/2})$. The other

part contains arbitrary $n^{1+\beta} - (n/2)^{3/2}$ edges (this is possible as long as $\beta < 1 - 3/\log n$).

Following a similar procedure, we can construct from the graphs of [46, 39] a graph with $n^{1+\beta}$ edges that requires normalized total share size $\Omega(n \log n)$.

7 Secret-sharing Schemes for k-Homogeneous Access Structures

We presented in Section 4 a connection between 2-server robust CDS protocols and (G, t)-secret-sharing schemes; this connection was used to construct better secret-sharing schemes for almost all graphs. We extend this connection to a more general case. We define (H, t)-secret-sharing schemes for hypergraphs H, and we present a connection between k-server robust CDS protocols and (H, t)-secret-sharing schemes. Then, this connection will be used to construct secret-sharing schemes for almost all k-homogeneous access structures. (Recall that an access structure is homogeneous if all its minimal authorized sets are of size k.)

7.1 Construction of (H, t)-Secret-Sharing Schemes

There is a one-to-one correspondence between uniform hypergraphs and homogeneous access structures; in particular, complete uniform hypergraphs correspond to threshold access structures. For a hypergraph $H = (V, E)$ and $t < |V|$, we define the access structure $\Gamma_{H,t}$ on V as $\Gamma_{H,t} = \Gamma_H \cup \Gamma_{t+1}$, where Γ_H is the access structure whose minimal authorized sets are the hyperedges in E, and Γ_t is the t-out-of-n threshold access structure.

Lemma 7.1. *Let $H = (V, E)$ be a k-hypergraph with $|V| = n$. Let $c(n, m)$ be the message size of k-server t-robust CDS protocols for functions with domain size n and secrets of size m. Then there is a (H, t)-secret-sharing scheme with maximum share size at most $k \cdot c(n, m) + \lceil \log(n) \rceil$.*

Proof. Let $f : [n]^k \to \{0, 1\}$ be the function defined as $f(i_1, \ldots, i_n) = 1$ if and only if $\{i_1, \ldots, i_n\} \in E$. Let $\mathcal{P}$ be a k-server t-robust CDS protocol with secrets in S for the function f, and let $Q_1, \ldots, Q_k$ be the servers of the protocol and let Π be a $(t + 1)$-out-of-n secret-sharing scheme.

Given a secret $s \in S$, we execute the protocol $\mathcal{P}$ for s, and, for every $1 \leq j \leq k$, we give the message of Q_j with input i to party $i \in V$. Moreover, we give to i a share of the scheme Π. The rest of the proof is analogous to that of Lemma 4.2. □

The following result is a direct consequence of Theorem 2.12 and Lemma 7.1.

Lemma 7.2. *Let $H = (V, E)$ be a k-hypergraph with $|V| = n$, and let $t \leq \min\{kn/2, 2^{\sqrt{n/k}}\}$ be an integer. Let $c(m, n)$ be the message size of k-server CDS protocols $\mathcal{P}$ with secrets of size m. Then, there is a (H, t)-secret-sharing scheme with maximum share size*

$$O\left(c(m, n) k^{3k} 2^k t^k \log^{2k-1} t \log^2 n\right).$$

If $\mathcal{P}$ is a linear protocol over $\mathbb{F}_{2^m}$, then the scheme is also linear. Furthermore, there is a (H, t)-secret-sharing scheme with secrets of size $\Theta(mtk^2 \log(t) \log^2(n))$ in which the normalized message size is

$$O\left(\frac{c(m, n)}{m} k^{3k-2} 2^k t^{k-1} \log^{2k-2} t\right).$$

7.2 Secret-Sharing Schemes for Almost All k-Hypergraphs

Following the same procedure we developed for graphs, we construct secret-sharing schemes for almost all k-hypergraphs from robust CDS protocols. In this case, we also use properties of the independent sets of almost all k-hypergraphs, shown in Lemma 7.3. Its proof is in Appendix A.

Lemma 7.3. *For every $k > 2$, for almost every k-hypergraph H the size of a maximum independent set in H is smaller than $k \log n$.*

Now we use Lemma 7.2 to construct secret-sharing schemes for almost all k-hypergraphs. In Theorem 7.4, we restrict our study to k-hypergraphs for a constant k. However, this approach is also valid for non-constant k.

Theorem 7.4. *Let $1 < k < n$ be a constant. Almost all k-homogeneous access structures with n parties can be realized by the following schemes.*

1. *A secret-sharing scheme with maximum share size $2^{\tilde{O}(\sqrt{k \log n})}$.*
2. *A linear secret-sharing scheme over $\mathbb{F}_2$ with maximum share size $\tilde{O}(n^{(k-1)/2})$.*
3. *A multi-linear secret-sharing scheme with normalized maximum share size $\tilde{O}(\log^{k-1} n)$ and exponentially long secrets of length 2^{n^k}.*

Proof. By Lemma 7.3, almost all k-hypergraphs have maximal independent sets of size at most $t = k \log n$. Hence, for almost every k-hypergraph H, any (H, t)-secret-sharing schemes realizes Γ_H. In order to construct these schemes we use Lemma 7.2 and different CDS protocols with different message size $c(m, n)$. Since k is constant, we are able to construct schemes with maximum share size

$$O\left(c(m,n) k^{3k} 2^k k^k \log^k n \log^{2k-1}(k \log n) \log^2(n)\right) = \tilde{O}\left(c(m,n) \log^{k+2} n\right),$$

and schemes with normalized maximum share size

$$O\left(\frac{c(m,n)}{m} k^{3k-2} 2^k k^{k-1} \log^{k-1} n \log^{2k-2}(k \log n)\right) = \tilde{O}\left(\frac{c(m,n)}{m} \log^{k-1} n\right).$$

First, we take the k-server CDS protocol from [61] that has message size $c(n, 1) = 2^{\tilde{O}(\sqrt{k \log n})}$. The resulting scheme has maximum share size $2^{\tilde{O}\sqrt{k \log n}}$. Second, we take the linear k-server CDS protocol from [19, 59], $c(n, 1) = n^{(k-1)/2}$. The resulting scheme has maximum share size $\tilde{O}(n^{(k-1)/2})$. Third, we take the multi-linear k-server CDS protocol from [2] in which the message size $c(n, m)/m = O(1)$, for exponentially long secrets of length 2^{n^k}. The resulting scheme has normalized maximum share size $\tilde{O}(\log^{k-1} n)$. □

Regarding lower bounds, it was proved in [4] that for almost all k-uniform access structures, the maximum share size for sharing a one-bit secret in a linear secret-sharing scheme is $\Omega(k^{-3/4} n^{-1/2} 2^{(h(k/n)/2)n})$, where $h(p)$ is the binary entropy function, namely, $h(p) = -p \log p - (1-p) \log(1-p)$.

References

[1] William Aiello, Yuval Ishai, and Omer Reingold. Priced oblivious transfer: How to sell digital goods. In Birgit Pfitzmann, editor, *EUROCRYPT 2001*, volume 2045 of *LNCS*, pages 118–134. Springer-Verlag, 2001.

[2] Benny Applebaum and Barak Arkis. On the power of amortization in secret sharing: d-uniform secret sharing and CDS with constant information rate. In Amos Beimel and Stefan Dziembowski, editors, *TCC 2018*, volume 11239 of *LNCS*, pages 317–344. Springer-Verlag, 2018.

[3] Benny Applebaum, Barak Arkis, Pavel Raykov, and Prashant Nalini Vasudevan. Conditional disclosure of secrets: Amplification, closure, amortization, lower-bounds, and separations. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017*, volume 10401 of *LNCS*, pages 727–757. Springer-Verlag, 2017.

[4] Benny Applebaum, Amos Beimel, Oriol Farràs, Oded Nir, and Naty Peter. Secret-sharing schemes for general and uniform access structures. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019*, volume 11478 of *LNCS*, pages 441–471. Springer-Verlag, 2019.

[5] Benny Applebaum, Amos Beimel, Oded Nir, and Naty Peter. Better secret-sharing via robust conditional disclosure of secrets. Cryptology ePrint Archive, Report 2020/080, 2020. `https://eprint.iacr.org/2020/080`.

[6] Benny Applebaum, Thomas Holenstein, Manoj Mishra, and Ofer Shayevitz. The communication complexity of private simultaneous messages, revisited. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018*, LNCS, pages 261–286. Springer-Verlag, 2018.

[7] Benny Applebaum and Prashant Nalini Vasudevan. Placing conditional disclosure of secrets in the communication complexity universe. In *10th ITCS*, pages 4:1–4:14, 2019.

[8] Nuttapong Attrapadung. Dual system encryption via doubly selective security: Framework, fully secure functional encryption for regular languages, and more. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 557–577. Springer-Verlag, 2014.

[9] László Babai, Anna Gál, and Avi Wigderson. Superpolynomial lower bounds for monotone span programs. *Combinatorica*, 19(3):301–319, 1999.

[10] Amos Beimel. Secret-sharing schemes: A survey. In Yeow Meng Chee, Zhenbo Guo, San Ling, Fengjing Shao, Yuansheng Tang, Huaxiong Wang, and Chaoping Xing, editors, *Coding and Cryptology – Third International Workshop, IWCC 2011*, volume 6639 of *LNCS*, pages 11–46. Springer-Verlag, 2011.

[11] Amos Beimel, Aner Ben-Efraim, Carles Padró, and Ilya Tyomkin. Multi-linear secret-sharing schemes. In Yehuda Lindell, editor, *TCC 2014*, volume 8349 of *LNCS*, pages 394–418. Springer-Verlag, 2014.

[12] Amos Beimel and Benny Chor. Universally ideal secret sharing schemes. *IEEE Trans. on Information Theory*, 40(3):786–794, 1994.

[13] Amos Beimel, Oriol Farràs, and Yuval Mintz. Secret-sharing schemes for very dense graphs. *J. of Cryptology*, 29(2):336–362, 2016.

[14] Amos Beimel, Oriol Farràs, Yuval Mintz, and Naty Peter. Linear secret-sharing schemes for forbidden graph access structures. Technical Report 2017/940, IACR Cryptology ePrint Archive, 2017. Full version of [15].

[15] Amos Beimel, Oriol Farràs, Yuval Mintz, and Naty Peter. Linear secret-sharing schemes for forbidden graph access structures. In Yael Kalai and Leonid Reyzin, editors, *TCC 2017*, volume 10678 of *LNCS*, pages 394–423. Springer-Verlag, 2017.

[16] Amos Beimel, Anna Gál, and Mike Paterson. Lower bounds for monotone span programs. *Computational Complexity*, 6(1):29–45, 1997.

[17] Amos Beimel, Yuval Ishai, Ranjit Kumaresan, and Eyal Kushilevitz. On the cryptographic complexity of the worst functions. In Yehuda Lindell, editor, *TCC 2014*, volume 8349 of *LNCS*, pages 317–342. Springer-Verlag, 2014.

[18] Amos Beimel, Eyal Kushilevitz, and Pnina Nissim. The complexity of multiparty PSM protocols and related models. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018*, volume 10821 of *LNCS*, pages 287–318. Springer-Verlag, 2018.

[19] Amos Beimel and Naty Peter. Optimal linear multiparty conditional disclosure of secrets protocols. In Thomas Peyrin and Steven Galbraith, editors, *ASIACRYPT 2018*, volume 11274 of *LNCS*, pages 332–362. Springer-Verlag, 2018.

[20] Amos Beimel and Enav Weinreb. Separating the power of monotone span programs over different fields. *SIAM J. on Computing*, 34(5):1196–1215, 2005.

[21] Michael Ben-Or, Shaffi Goldwasser, and Avi Wigderson. Completeness theorems for noncryptographic fault-tolerant distributed computations. In *20th STOC*, pages 1–10, 1988.

[22] Josh Cohen Benaloh and Jerry Leichter. Generalized secret sharing and monotone functions. In Shaffi Goldwasser, editor, *CRYPTO '88*, volume 403 of *LNCS*, pages 27–35. Springer-Verlag, 1988.

[23] Michael Bertilsson and Ingemar Ingemarsson. A construction of practical secret sharing schemes using linear block codes. In Jennifer Seberry and Yuliang Zheng, editors, *AUSCRYPT '92*, volume 718 of *LNCS*, pages 67–79. Springer-Verlag, 1992.

[24] George Robert Blakley. Safeguarding cryptographic keys. In *Proc. of the 1979 AFIPS National Computer Conference*, volume 48 of *AFIPS Conference proceedings*, pages 313–317. AFIPS Press, 1979.

[25] Carlo Blundo, Alfredo De Santis, Roberto De Simone, and Ugo Vaccaro. Tight bounds on the information rate of secret sharing schemes. *Designs, Codes and Cryptography*, 11(2):107–122, 1997.

[26] Carlo Blundo, Alfredo De Santis, Luisa Gargano, and Ugo Vaccaro. On the information rate of secret sharing schemes. *Theoretical Computer Science*, 154(2):283–306, 1996.

[27] Carlo Blundo, Alfredo De Santis, Antonio Giorgio Gaggia, and Ugo Vaccaro. New bounds on the information rate of secret sharing schemes. *IEEE Trans. on Information Theory*, 41(2):549–553, 1995.

[28] Béla Bollobás. *Random graphs*. Cambridge, 2nd edition edition, 2001.

[29] Béla Bollobás and Andrew Thomason. Graphs which contain all small graphs. *European Journal of Combinatorics*, 2(1):13–15, 1981.

[30] Ernest F. Brickell. Some ideal secret sharing schemes. *Journal of Combin. Math. and Combin. Comput.*, 6:105–113, 1989.

[31] Ernest F. Brickell and Daniel M. Davenport. On the classification of ideal secret sharing schemes. *J. of Cryptology*, 4(73):123–134, 1991.

[32] Ernest F. Brickell and Douglas R. Stinson. Some improved bounds on the information rate of perfect secret sharing schemes. *J. of Cryptology*, 5(3):153–166, 1992.

[33] Renato M. Capocelli, Alfredo De Santis, Luisa Gargano, and Ugo Vaccaro. On the size of shares for secret sharing schemes. *J. of Cryptology*, 6(3):157–168, 1993.

[34] David Chaum, Claude Crépeau, and Ivan Damgård. Multiparty unconditionally secure protocols. In *20th STOC*, pages 11–19, 1988.

[35] Benny Chor and Eyal Kushilevitz. Secret sharing over infinite domains. *J. of Cryptology*, 6(2):87–96, 1993.

[36] Ronald Cramer, Ivan Damgård, and Ueli Maurer. General secure multi-party computation from any linear secret-sharing scheme. In Bart Preneel, editor, *EUROCRYPT 2000*, volume 1807 of *LNCS*, pages 316–334. Springer-Verlag, 2000.

[37] László Csirmaz. The dealer's random bits in perfect secret sharing schemes. *Studia Sci. Math. Hungar.*, 32(3–4):429–437, 1996.

[38] László Csirmaz. The size of a share must be large. *J. of Cryptology*, 10(4):223–231, 1997.

[39] László Csirmaz. Secret sharing schemes on graphs. Technical Report 2005/059, Cryptology ePrint Archive, 2005. `eprint.iacr.org/`.

[40] László Csirmaz. An impossibility result on graph secret sharing. *Des. Codes Cryptography*, 53(3):195–209, 2009.

[41] László Csirmaz. Secret sharing on the d-dimensional cube. *Designs, Codes and Cryptography*, 74(3):719–729, 2015.

[42] László Csirmaz and Péter Ligeti. Secret sharing on large girth graphs. *Cryptogr. Commun.*, 11(3):399–410, 2019.

[43] László Csirmaz and Gábor Tardos. Optimal information rate of secret sharing schemes on trees. *IEEE Trans. Inf. Theory*, 59(4):2527–2530, 2013.

[44] Yvo Desmedt and Yair Frankel. Shared generation of authenticators and signatures. In Joan Feigenbaum, editor, *CRYPTO '91*, volume 576 of *LNCS*, pages 457–469. Springer-Verlag, 1991.

[45] Marten van Dijk. On the information rate of perfect secret sharing schemes. *Des. Codes Cryptography*, 6(2):143–169, 1995.

[46] Marten van Dijk. On the information rate of perfect secret sharing schemes. *Designs, Codes and Cryptography*, 6:143–169, 1995.

[47] Paul Erdös and László Pyber. Covering a graph by complete bipartite graphs. *Discrete Mathematics*, 170(1–3):249–251, 1997.

[48] Paul Erdös and Alfréd Rényi. On random graphs. I. *Publ. Math. Debrecen*, 6:290–297, 1959.

[49] Oriol Farràs, Tarik Kaced, Sebastià Martín, and Carles Padró. Improving the linear programming technique in the search for lower bounds in secret sharing. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018*, LNCS, pages 597–621. Springer-Verlag, 2018.

[50] Anna Gál and Pavel Pudlák. Monotone complexity and the rank of matrices. *Inform. Process. Lett.*, 87:321–326, 2003.

[51] Romain Gay, Iordanis Kerenidis, and Hoeteck Wee. Communication complexity of conditional disclosure of secrets and attribute-based encryption. In Rosario Gennaro and Matthew Robshaw, editors, *CRYPTO 2015*, volume 9216 of *LNCS*, pages 485–502. Springer-Verlag, 2015.

[52] Yael Gertner, Yuval Ishai, Eyal Kushilevitz, and Tal Malkin. Protecting data privacy in private information retrieval schemes. *J. of Computer and System Sciences*, 60(3):592–629, 2000.

[53] Viput Goyal, Omkant Pandey, Amit Sahai, and Brent Waters. Attribute-based encryption for fine-grained access control of encrypted data. In *13th CCS*, pages 89–98, 2006.

[54] Geoffrey R. Grimmett and Colin J. H. McDiarmid. On colouring random graphs. *Mathematical Proceedings of the Cambridge Philosophical Society*, 77(2):313324, 1975.

[55] Mitsuru Ito, Akira Saito, and Takao Nishizeki. Secret sharing schemes realizing general access structure. In *Globecom 87*, pages 99–102, 1987. Journal version: Multiple assignment scheme for sharing secret. *J. of Cryptology*, 6(1), 15-20, 1993.

[56] Mauricio Karchmer and Avi Wigderson. On span programs. In *8th Structure in Complexity Theory*, pages 102–111, 1993.

[57] Ehud D. Karnin, Jonathan W. Greene, and Martin E. Hellman. On secret sharing systems. *IEEE Trans. on Information Theory*, 29(1):35–41, 1983.

[58] A. D. Korshunov. On the number of monotone Boolean functions. *Probl. Kibern*, 38:5–108, 1981.

[59] Tianren Liu and Vinod Vaikuntanathan. Breaking the circuit-size barrier in secret sharing. In *50th STOC*, pages 699–708, 2018.

[60] Tianren Liu, Vinod Vaikuntanathan, and Hoeteck Wee. Conditional disclosure of secrets via non-linear reconstruction. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017*, volume 10401 of *LNCS*, pages 758–790. Springer-Verlag, 2017.

[61] Tianren Liu, Vinod Vaikuntanathan, and Hoeteck Wee. Towards breaking the exponential barrier for general secret sharing. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018*, volume 10820 of *LNCS*, pages 567–596. Springer-Verlag, 2018.

[62] J. Martí-Farré and C. Padró. On secret sharing schemes, matroids and polymatroids. *J. Mathematical Cryptology*, 4(2):95–120, 2010.

[63] Yuval Mintz. Information ratios of graph secret-sharing schemes. Master's thesis, Dept. of Computer Science, Ben Gurion University, 2012.

[64] Moni Naor and Avishai Wool. Access control and signatures via quorum secret sharing. In *3rd CCS*, pages 157–167, 1996.

[65] Peter Nelson. Almost all matroids are nonrepresentable. *Bulletin of the London Mathematical Society*, 50(2):245–248, 2018.

[66] Naty Peter. *Secret-Sharing Schemes and Conditional Disclosure of Secrets Protocols*. PhD thesis, Ben-Gurion University of the Negev, 2020.

[67] Toniann Pitassi and Robert Robere. Lifting nullstellensatz to monotone span programs over any field. In *50th STOC*, pages 1207–1219, 2018.

[68] Lajos Rónyai, László Babai, and Murali K. Ganapathy. On the number of zero-patterns of a sequence of polynomials. *Journal of the AMS*, 14(3):717–735, 2001.

[69] Adi Shamir. How to share a secret. *Communications of the ACM*, 22:612–613, 1979.

[70] Claude E. Shannon. The synthesis of two-terminal switching circuits. *Bell Syst. Tech. J.*, 28(1):59–98, 1949.

[71] Gustavus J. Simmons. How to (really) share a secret. In Shaffi Goldwasser, editor, *CRYPTO '88*, volume 403 of *LNCS*, pages 390–448. Springer-Verlag, 1990.

[72] Hung-Min Sun and Shiuh-Pyng Shieh. Secret sharing in graph-based prohibited structures. In *INFOCOM '97*, pages 718–724, 1997.

[73] Brent Waters. Ciphertext-policy attribute-based encryption: An expressive, efficient, and provably secure realization. In Dario Catalano, Nelly Fazio, Rosario Gennaro, and Antonio Nicolosi, editors, *PKC 2011*, volume 6571 of *LNCS*, pages 53–70. Springer-Verlag, 2011.

[74] Hoeteck Wee. Dual system encryption via predicate encodings. In Yehuda Lindell, editor, *TCC 2014*, volume 8349 of *LNCS*, pages 616–637. Springer-Verlag, 2014.

[75] Ingo Wegener. *The Complexity of Boolean Functions*. Wiley-Teubner Series in Computer Science. B. G. Teubner and John Wiley, 1987.

A Maximal independent sets

In this appendix, we present the proofs of Lemma 2.14, Lemma 2.15 and Lemma 7.3.

Proof of Lemma 2.14. Let $1/n \leq p \leq 1/2$. In this case, we can write $p = n^{-\alpha}$ for some $1/\log n \leq \alpha \leq 1$. The probability that an edge e is not to be in G is $1-p$. The probability that any t vertices $v_1, \ldots, v_t$ are an independent set in G is

$$p' = (1-p)^{\binom{t}{2}}.$$

Let p'' the probability that there is an independent set of size t in G. By the union bound,

$$p'' \leq \binom{n}{t} p' \leq \left(\frac{en}{t}\right)^t (1-p)^{\binom{t}{2}} \leq 2^r, \text{ where}$$

$$r = t\left(\log e + \log n - \log t + \frac{1}{2}(t-1)\log(1-p)\right). \tag{1}$$

Let $t = 4\log n/\log(\frac{1}{1-p})$. For this value of t, it holds $\frac{1}{2}(t-1)\log(1-p) = -2\log n + \frac{1}{2}\log\left(\frac{1}{1-p}\right)$, and so

$$\begin{aligned} r &= t\left(\log e + \log n - \log t - 2\log n + \frac{1}{2}\log(\frac{1}{1-p})\right) \\ &= t\left(-\log n - \log t + \frac{1}{2}\log(\frac{1}{1-p}) + \log e\right). \end{aligned}$$

Since $1/\log n \leq \alpha \leq 1$, we have $\frac{n}{n-1} \leq \frac{1}{1-p} \leq 2$ and $0 < \log\frac{1}{1-p} \leq 1$. Hence, when $n \to \infty$, we have that $t \to \infty$ and $r \to -\infty$. Therefore, the probability that there is an independent set of size t in G tends to 0 when $n \to \infty$.

Using that $\log(1-x) = \sum_{i=1}^{\infty} -\frac{x^i}{i} < -x$ for any $0 < x < 1$, we get that

$$t = \frac{4\log n}{\log\left(\frac{1}{1-n^{-\alpha}}\right)} = \frac{4\log n}{-\log(1-n^{-\alpha})} < 4n^{\alpha}\log n.$$

Now let $p = 1 - n^{-\alpha}$. Following the same argument, we get that the probability that there is an independent set of size t in G is

$$p'' \leq \binom{n}{t} p' \leq \left(\frac{en}{t}\right)^t n^{-\alpha\binom{t}{2}} = \left(\frac{e}{t}\right)^t n^{t\left(1-\alpha\frac{t-1}{2}\right)}.$$

If $\alpha\frac{t-1}{2} - 1 \leq \log^{-1/2} n$, then $p'' \to 0$ when $n \to \infty$. □

Proof of Lemma 2.15. In $\mathscr{G}(n,p)$, the expected value of the degree of each edge is smaller than $pn = \mu = \omega(\log n)$. By a Chernov bound, for any $\delta \geq 0$, the probability that the degree of one vertex is at most $(1+\delta)\mu$, is at most $(\frac{e^{\delta}}{(1+\delta)^{1+\delta}})^{\mu}$. We take $\delta = 1$. By the union bound, the probability that at least one vertex has degree greater than 2μ is at most

$$n\left(\frac{e^{\delta}}{(1+\delta)^{1+\delta}}\right)^{\mu} = n\left(\frac{e}{4}\right)^{\omega(\log n)} = n^{1-\omega(1)} = n^{-\omega(1)}.$$

Hence, with probability at least $1 - n^{-\omega(1)}$ a graph in $\mathscr{G}(n,p)$ has degree at most 2μ. □

Proof of Lemma 7.3. We construct a random k-hypergraph H taking subsets of size k with probability $p = 1/2$. The probability that $t = k\log n$ vertices $v_1, \ldots, v_t$ contains all $\binom{t}{k}$ subsets of size k is $p' = p^{\binom{t}{k}}$, and the probability that it happens for a set of size t in H is less than

$$\binom{n}{t} p' \leq \frac{\binom{n}{t}}{2^{\binom{t}{k}}} =\leq 2^{t\log\left(\frac{en}{t}\right) - \left(\frac{t}{k}\right)^k} \leq 2^{k^2\log^2 n - \log^k n},$$

which tends to $2^{-\log^k n}$ when $n \to \infty$ for $k > 2$. The probability that there is no independent set of size t is greater than $1 - \binom{n}{t}p'$, which tends to 1 when $n \to \infty$. Therefore, almost all k-hypergraphs have independent sets of size smaller than $k\log n$. □