CONJURA · OPEN PROBLEM

Beating $\log n$ *for Computational* 2*-out-of-*n *Secret Sharing*

Share size $\delta \log n$ *for some* $\delta < 1$*, with public information*

Status. Statement: AI-written from Damiano Abram, Amos Beimel and Yuval Ishai, *Cryptography from Planted Graphs: Security with Logarithmic-Size Messages*, IACR ePrint 2023/1929. Not yet checked by a human. The information-theoretic share size is exactly $\log n$; Shamir's scheme matches it computationally. The source proves a $\frac{1}{5} \log \log n$ lower bound for the computational setting with public information, leaving a $\log n$ versus $\log \log n$ gap, and proves that closing it upward is equivalent to a planted-subgraph problem.

Category. *Secret Sharing.*

Conjecture (informal). *Sharing one bit among* n *parties so that any two of them can recover it takes* $\log n$ *bits per share if the scheme must be perfectly private, and Shamir's scheme achieves that. Allowing computational security, and allowing the dealer to publish some public information alongside the shares, opens the possibility of doing better — and unusually for secret sharing, there is no obvious lower bound at all in that setting. The source narrows the range to between* $\frac{1}{5} \log \log n$ *and* $\log n$ *and then does something more useful than either bound: it proves that beating* $\log n$ *by any constant factor is* equivalent *to a concrete planted-graph problem, hiding a large clique and a large independent set in the same graph so that a random node of one cannot be told from a random node of the other.*

1 The setting

In secret sharing with public information the dealer publishes a string I in addition to handing each party its share; correctness and privacy are as usual, with I available to everyone including the adversary. The model is what makes sub-logarithmic share size conceivable: for 2-out-of-n threshold sharing of a one-bit secret, the information-theoretic share size is $\log n$ and no better, and Shamir's

scheme meets it, but as the source observes, "in the context of secret-sharing schemes with public information, there is no obvious lower bound on the share size."

The source's contribution here is a translation. A 2-out-of-n scheme with public information and share size ℓ is equivalent to a multipartite planted-clique problem: from I one derives an n-partite graph with 2^ℓ nodes per part, the nodes of part i being the possible shares of party P_i, with an edge between every pair of shares that reconstructs to 1. Correctness then implies that when the secret is 1 the graph hides an n-node clique, and when it is 0 an n-node independent set, one node per part in either case; privacy says the two distributions are indistinguishable even given one node of the hidden subgraph. Beating $\log n$ means hiding an n-node clique in a graph with fewer than n^2 nodes, and the source reports that in that parameter regime the attacks of Kučera and of Alon, Krivelevich and Sudakov recover the clique for every graph distribution it tried, with multipartiteness making the goal look harder still.

A random-partitioning argument then removes the multipartite structure and yields the clean equivalence reproduced as Theorem 1 below. In the other direction the source proves a lower bound: share size at least $\frac{1}{5}\log\log n$, by the observation that a 2-out-of-n scheme induces a 2-out-of-n' scheme for every $n' \le n$ whose security does not degrade while the public information shrinks.

So the range is $[\frac{1}{5}\log\log n,\ \log n]$, and the source asks: "What is the minimal share size of computationally secure 2-out-of-n secret sharing with public information? Is it possible to beat the information-theoretic $\log n$ bound, even by a constant factor?"

2 Notation and parameters

- n is the number of parties; the access structure is 2-out-of-n and the secret is one bit.
- ℓ is the share size in bits; I is the public information, available to all parties and to the adversary.
- $\delta \in (0,1)$ is the constant by which the conjecture asks to beat $\log n$, i.e. the target share size is $\delta \log n$.
- In the graph formulation, N is the number of nodes of the graph, $\beta \in (1/2,1)$ the exponent giving the planted subgraphs N^β nodes each.

- Security is computational: against non-uniform $n^{o(\log n)}$-time adversaries with distinguishing advantage $n^{-\Omega(1)}$, which is the regime the source's own positive results achieve elsewhere and the regime its table reports.

3 The conjecture

Definition 1 (2-out-of-n secret sharing with public information). A dealer holding a bit b outputs public information I and shares $s_1, \ldots, s_n \in \{0,1\}^\ell$. *Correctness*: for every $i \neq j$ there is a reconstruction procedure recovering b from (I, s_i, s_j). *Privacy*: no admissible adversary given I and any single share s_i can distinguish $b = 0$ from $b = 1$ better than the allowed advantage. The share size is ℓ; the size of I is not counted.

Conjecture 1 (Beating $\log n$). *There is a constant $0 < \delta < 1$ and a computational 2-out-of-n secret-sharing scheme with public information (Definition 1) whose share size is $\delta \cdot \log n$.*

Theorem 1 (Equivalent planted-subgraph form, [ABI23, Theorem 2.6, informal]). *The following are equivalent.*

- *There exists a constant $0 < \delta < 1$ for which there is a computational 2-out-of-n secret-sharing scheme with public information and $\delta \cdot \log n$ share size.*
- *There exist a constant $1/2 < \beta < 1$ and a distribution $\mathcal{D}$ of triples (G, C, J), where G is an N-node graph, C an N^β-node clique in G and J an N^β-node independent set in G, such that it is hard to distinguish (G, c) from (G, j) for c and j random nodes of C and J respectively.*

The source writes I for the independent set here and also for the public information of Definition 1; the independent set is renamed J throughout this statement to keep the two apart.

Remark 1 (Both directions are open, and the gap is large). The source poses this as a question and predicts nothing. The known range is $\frac{1}{5} \log\log n \leq \ell \leq \log n$, so a resolution can come from either end: a scheme at $\delta \log n$ for some $\delta < 1$, or a lower bound above $\Omega(\log\log n)$ — ideally $\Omega(\log n)$, which would show that computational security and public information buy nothing here

at all. Note that Theorem 1 makes a negative answer a statement about graphs too: ruling out the scheme is the same as ruling out the distribution $\mathcal{D}$.

Remark 2 (Why this is not a standard planted-clique problem). The source flags the difference, and it matters for anyone attacking Theorem 1's second item. In traditional planted problems the planting guarantees the object occurs essentially once, so there is a natural search version: find it. Here $\mathcal{D}$ may be such that every node lies in many planted cliques and independent sets — indeed for the multipartite version there is an n-partite graph with only 4 nodes per part in which every node lies in both a clique and an independent set of size n, by Beimel and Franklin — so no natural search problem presents itself. What is more, the information-theoretic impossibility of beating $\log n$ implies the decision problem above *is* solvable by a computationally unbounded distinguisher. So the target is a genuinely computational gap, not an information-theoretic one.

Remark 3 (The evidence against, such as it is). Beating $\log n$ requires hiding an n-node clique in a graph on fewer than n^2 nodes, and the source reports that in that regime the classical clique-recovery attacks succeeded on every distribution it tried. That is evidence, not a proof, and the source presents it as such. It is the main reason to treat a positive answer as the surprising direction.

Remark 4 (Neighbouring questions in the same list). The source's open-questions section also asks whether the limitations of its positive results can be removed — security only against $n^{o(\log n)}$-time adversaries, and distinguishing advantage only $n^{-\Omega(1)}$ rather than negligible, where making it negligible costs a super-constant multiplicative overhead — and whether computational-statistical gaps for secret sharing with 1-bit shares extend to domain size 3. Neither is this statement.

4 Bibliography

[ABI23] Damiano Abram, Amos Beimel and Yuval Ishai. *Cryptography from planted graphs: security with logarithmic-size messages.* IACR ePrint 2023/1929. The source. The open question is in Section 1.1.5, page

8; the graph translation and Theorem 2.6 are in Section 2.4, pages 16–17; the $\frac{1}{5} \log \log n$ lower bound is Theorem 8.1.

[Sha79] Adi Shamir. *How to share a secret.* Communications of the ACM, 22(11):612–613, November 1979. The scheme achieving $\log n$ share size, which the conjecture asks to beat.

[KN90] Joe Kilian and Noam Nisan. Private communication, 1990. One of the two sources of the information-theoretic $\log n$ lower bound.

[CCX13] Ignacio Cascudo, Ronald Cramer and Chaoping Xing. *Bounds on the threshold gap in secret sharing and its applications.* IEEE Transactions on Information Theory, 2013. The other source of that bound.

[BF07] Amos Beimel and Matthew K. Franklin. *Weakly-private secret sharing schemes.* 4th Theory of Cryptography Conference (TCC), 2007. The source of the 4-nodes-per-part example showing every node can lie in both a clique and an independent set of size n.

[Kuc95] Luděk Kučera. *Expected complexity of graph partitioning problems.* Discrete Applied Mathematics, 1995. One of the two clique-recovery attacks the source reports succeed in the sub-n^2 regime.

[AKS98] Noga Alon, Michael Krivelevich and Benny Sudakov. *Finding a large hidden clique in a random graph.* Random Structures and Algorithms 13, 1998. The other attack the source reports.