Weighted Secret Sharing from Wiretap Channels

Fabrice Benhamouda[*] Shai Halevi[*] Lev Stambler[†]

February 8, 2023

Abstract

Secret-sharing allows splitting a piece of secret information among a group of shareholders, so that it takes a large enough subset of them to recover it. In *weighted* secret-sharing, each shareholder has an integer weight, and it takes a subset of large-enough weight to recover the secret. Schemes in the literature for weighted threshold secret sharing either have share sizes that grow linearly with the total weight, or ones that depend on huge public information (essentially a garbled circuit) of size (quasi)polynomial in the number of parties.

To do better, we investigate a relaxation, (α, β)-ramp weighted secret sharing, where subsets of weight βW can recover the secret (with W the total weight), but subsets of weight αW or less cannot learn anything about it. These can be constructed from standard secret-sharing schemes, but known constructions require long shares even for short secrets, achieving share sizes of $\max\big(W, \frac{|\text{secret}|}{\epsilon}\big)$, where $\epsilon = \beta - \alpha$. In this note we first observe that simple rounding let us replace the total weight W by N/ϵ, where N is the number of parties. Combined with known constructions, this yields share sizes of $O\big(\max(N, |\text{secret}|)/\epsilon\big)$.

Our main contribution is a novel connection between weighted secret sharing and wiretap channels, that improves or even eliminates the dependence on N, at a price of increased dependence on $1/\epsilon$. We observe that for certain additive-noise $(\mathcal{R}, \mathcal{A})$ wiretap channels, any semantically secure scheme can be naturally transformed into an (α, β)-ramp weighted secret-sharing, where α, β are essentially the respective capacities of the channels $\mathcal{A}, \mathcal{R}$. We present two instantiations of this type of construction, one using Binary Symmetric wiretap Channels, and the other using additive Gaussian Wiretap Channels. Depending on the parameters of the underlying wiretap channels, this gives rise to (α, β)-ramp schemes with share sizes $|\text{secret}|/\text{poly}(\epsilon \log N)$ or even just $|\text{secret}|/\text{poly}(\epsilon)$.

1 Introduction

Secret sharing [Sha79, Bla79] allows a dealer to split some secret information among multiple parties, giving each party an individual share, so that large enough subsets of shareholder can recover the secret, but small subsets cannot learn any partial information about it. Such schemes are typically parametrized by the number of parties N and a threshold $T \leq N$, such that it takes at least T parties to recover the secret.

Weighted secret sharing (WSS) is similar, except that each shareholder j has an integer weight w_j, it takes a "heavy enough" subsets to recover the secret, while "light" subsets cannot learn any partial

[*] Algorand Foundation, USA

[†] Work done while in Algorand Foundation, USA

information about it. The threshold $T \in [N]$ is replaced by $\tau \in (0,1)$, such that it takes shareholders of aggregate weight τW to recover the secret (where W is the total weight, $W = \sum_{j\in[N]} w_j$).

One method of implementing WSS is to rely on standard secret-sharing with $N' = W$ and $T' = \tau W$, giving w_j shares to a shareholder j with weight w_j. While this solution can achieve good rate for long secrets (see section 3.1), it is very wasteful for short ones, as its share sizes grow linearly with the weight. Prior work on weighted secret sharing explored other solutions (e.g., using Chinese remaindering) or limited models (e.g., specific weight hierarchies). But they all still feature either linear dependency of the share-size on W, severe restrictions to the access structures that can be realized, or huge public information that must be broadcasted to everyone alongside the individual shares. (See more discussion in section 1.2 below.)

In an attempt to do better, in this work we consider the relaxed model of ramp secret-sharing [BM84], that has a fuzzy threshold. Specifically, an (α, β)-ramp weighted secret sharing scheme allows any subset of aggregate weight at least βW to recover the secret, but subsets of weight αW or less cannot learn any information about it. Such gaps were considered often in the literature for standard secret-sharing schemes, but to our knowledge were not studied in the context of weighted secret sharing.

It is not hard to see (and we describe it explicitly in section 3) that this relaxation enables shorter secrets, just by keeping only a $1/\epsilon$ precision for the weights, where $\epsilon = \beta - \alpha$. Rather than linear dependence on the weights, we now get linear dependence on N/ϵ (where the dependence on the number of parties N is due to the accumulation of rounding errors in this limited-precision approximation).

Beyond this simple observation, the main technical meat in this work is a novel blueprint for (α, β)-ramp WSS schemes, by exploring a surprising connection to secure transmission schemes for wiretap channels. These constructions reduce or even eliminate the dependence on N, at the price of potentially worse (but still polynomial) dependence on $1/\epsilon$. We note that the field of wiretap coding is an ongoing line of research with an aim of decreasing dependence on $1/\epsilon$. Any advances in wiretap coding can easily be applied to WSS with our construction.

1.1 Overview of Our Techniques

The starting point for our new blueprint is the following approach: On input $\boldsymbol{s}$, the dealer gives each shareholder j a noise vector $\boldsymbol{e}_j$, whose magnitude depends on their weight, and publishes the value $\boldsymbol{g} = \mathsf{Enc}(\boldsymbol{s}) + \sum_j \boldsymbol{e}_j$, where $\mathsf{Enc}(\cdot)$ is some encoding function. Given the public $\boldsymbol{g}$ and their individual $\boldsymbol{e}_j$'s, the only information that a set T of shareholder has on the secret $\boldsymbol{s}$ is the value

$$\boldsymbol{g}_T = \boldsymbol{g} - \sum_{j\in T} \boldsymbol{e}_j = \mathsf{Enc}(\boldsymbol{s}) + \sum_{j\notin T} \boldsymbol{e}_j.$$

We can therefore associate with each subset T an additive-noise channel $C_T : x \mapsto x + \sum_{j\notin T} \boldsymbol{e}_j$, such that the information that T learns about $\boldsymbol{s}$ is exactly the received value $C_T(\mathsf{Enc}(\boldsymbol{s}))$. We are seeking an encoding function $\mathsf{Enc}(\cdot)$ so that:

- Any qualified set S can recover $\boldsymbol{s}$ from $C_S(\mathsf{Enc}(\boldsymbol{s}))$;
- For any unqualified set T, seeing $C_T(\mathsf{Enc}(\boldsymbol{s}))$ yields no information on $\boldsymbol{s}$.

Intuitively, the smaller (or "lighter") the set is, the more error components it is missing, so the more noisy its channel will be. Consider now $\mathcal{R}$ which is "the most noisy channel" for any qualified set, and $\mathcal{A}$ which is "the least noisy channel" for any unqualified set. We can hope that $\mathcal{R}$ is less noisy than $\mathcal{A}$, and use a good transmission scheme for the wiretap channel $(\mathcal{R}, \mathcal{A})$, with receiver channel $\mathcal{R}$ and adversary channel $\mathcal{A}$.

Trying to flesh out this approach, we need to associate an error distribution $\mathcal{D}_{w_j}$ to every weight $w_j \in \mathbb{N}$, so that whenever $\sum_{j\in A} w_j > \sum_{j\in B} w_j$ it holds that $\sum_{j\in A} \mathcal{D}_{w_j}$ is "more error" than $\sum_{j\in B} \mathcal{D}_{w_j}$. Then we need to find two concrete channels $\mathcal{R}, \mathcal{A}$ such that

- $\mathcal{R}$ is at least as noisy as C_Q for any qualified set Q with weight $\geq \beta W$.
- $\mathcal{A}$ is at most as noisy as C_U for any unqualified set U with weight $\leq \alpha W$.

If $\mathcal{R}$ is less noisy than $\mathcal{A}$, then we can use a good transmission scheme for the wiretap channel $(\mathcal{R}, \mathcal{A})$ to implement our (α, β)-ramp WSS scheme. The parameters of this WSS scheme can be derived from those of the underlying wiretap scheme.

1.1.1 Binary Symmetric Channels

Trying to instantiate this approach with binary symmetric channels, we associate with each weight w_j an error probability p_j and the corresponding Bernoulli random variable

$$\mathcal{D}_j = \begin{cases} 1 & \text{with probability } p_j \\ 0 & \text{with probability } 1 - p_j. \end{cases}$$

One problem that we face here is that the error probability does not add up linearly. If we set (say) $p_j = w_j/W$, it is not hard to find instances where $\sum_{j\in A} w_j > \sum_{j\in B} w_j$ and yet $\sum_{i\in A} \mathcal{D}_j \bmod 2$ has smaller error probability than $\sum_{j\in B} \mathcal{D}_j \bmod 2$, as the following example shows.

A problematic example. Consider three parties with $w_1 = w_2 = 13$ and $w_3 = 24$, so $W = 50$ and we have $\Pr[\mathcal{D}_1 = 1] = \Pr[\mathcal{D}_2 = 1] = 13/50 = 0.26$ and $\Pr[\mathcal{D}_3 = 1] = 24/50 = 0.48$. Let $A = \{1, 2\}$ and $B = \{3\}$, so the aggregate weight of A is 26, larger that the weight of B which is 24. On the other hand, we have

$$\Pr[\mathcal{D}_1 \oplus \mathcal{D}_2 = 1] = 0.26 + 0.26 - 0.26^2 = 0.4525 < 0.48 = \Pr[\mathcal{D}_3 = 1],$$

so the error rate for A is *lower* that for B.

Clearly, the reason for this example is the cancellation due to the term 0.26^2. This cancellation effect can be reduced by scaling down the probabilities. For example, if we set $p_j = w_j/2W$ rather than $p_j = w_j/W$, then we get $\Pr[\mathcal{D}_1 = 1] = \Pr[\mathcal{D}_2 = 1] = 0.13$ and $\Pr[\mathcal{D}_3 = 1] = 0.24$, and therefore

$$Pr[\mathcal{D}_1 \oplus \mathcal{D}_2 = 1] = 0.13 + 0.13 - 0.13^2 = 0.2431 > 0.24 = \Pr[\mathcal{D}_3 = 1].$$

Our construction. The BSC-based construction that we present in section 5 comes with a scaling parameter γ (that depends on α, β) that controls the noise level. Namely, a weight w shareholder gets an error vector in which every bit is one with probability $\gamma \cdot w/W$. Setting γ small enough ensures that any subset of weight $\leq (1 - \beta)W$ will have error rate smaller than every subset of aggregate weight $\geq (1 - \alpha)W$. This allows us find a good wiretap channel scheme for the resulting channels, and therefore construct an (α, β)-ramp WSS scheme.

A drawback of scaling down, however, is that the error rates of $\mathcal{R}$ and $\mathcal{A}$ become quite close, of distance only $O(\epsilon^2)$ (where $\epsilon = \beta - \alpha$). We therefore need to use codes of fairly large block-length, making the share-size grow polynomially in $1/\epsilon$, for a large polynomial. The details are described in section 5.

1.1.2 Additive Gaussian Channels

Trying to improve over the BSC-based construction, we turn our attention to additive white Gaussian noise (AWGN) channels. These channels have the advantage that their noise is additive: adding Gaussian variables with variance σ_1^2 and σ_2^2 yields another Gaussian with variance $\sigma_1^2 + \sigma_2^2$. This lets us avoid the scaling-down issue, potentially yielding better parameters.

We thus associate each weight, $w \in \mathbb{N}$ with the Normal random variable $\mathcal{N}(0, w/W)$, i.e. zero-mean with variance w/W (stdev $= \sqrt{w/W}$). Due to additivity, the aggregate random variable for a set A is itself a Normal variable,

$$\sum_{j \in A} \mathcal{N}(0, w_j/W) = \mathcal{N}(0, \sum_{j \in A} w_j/W).$$

This implies that whenever S has higher weight than T, the channel C_T has more error than the channel C_S.

For any $\beta > \alpha$, we can therefore construct an (α, β)-ramp WSS scheme from a good transmission scheme for the AGWN wiretap channel $(\mathcal{R}, \mathcal{A})$, where

$$\mathcal{R} : x \mapsto x + \mathcal{N}(0, 1-\beta) \;\text{ and }\; \mathcal{A} : x \mapsto x + \mathcal{N}(0, 1-\alpha).$$

Indeed, since $\beta > \alpha$ then $\mathcal{A}$ is more noisy than $\mathcal{R}$.

Quantization. Recall that in our approach the dealer draws the share for each party j as $\boldsymbol{e}_g \leftarrow \mathcal{N}(0, w_j/W)$, and in addition it publishes $\boldsymbol{g} = \mathsf{Enc}(\boldsymbol{s}) + \sum_{j \in [N]} \boldsymbol{e}_j$. But the noise entries are real random variables, so we need to quantize them in order to use them in the scheme, and the required precision for this quantization factors into the share sizes. The channels that we deal with are therefore not exactly AWGN but rather a finite-precision approximation, and we need to show that the approximation is good enough to maintain correctness and secrecy.

It is easy to see that secrecy is maintained, no matter how these real variables are quantized. Indeed, the underlying wiretap channel ensures secrecy even against adversaries that get an infinite precision random variables, and any quantized version will just be a processing of those real variables.

For correctness, we note that rounding cannot move the integers that the decoder sees by too much. Therefore, the correct codeword cannot be too much farther away than in the infinite-precision case, and the decoder will therefore succeed with almost the same probability.

1.2 Prior Work

Ramp secret-sharing (without weights) was introduced by Blakley and Meadows [BM84]. A textbook construction for a ramp-scheme with good rate based on standard "packed secret sharing" can be found, e.g., in [CDN15, 11.4.2] (and is described in section 3.1 below).

Some early work on weighted secret sharing was cast against the backdrop of general access structures. Beimel et al. [BTW05] characterized the weighted (strict) thresholds access structures that admit *ideal* schemes, where the share size is equal to the secret size, proving that only few specific threshold structures can be realized this way.

Beimel and Weinreb [BW06] showed that any threshold access structure can be realized using shares of size quasiPoly($N \log W$) times the secret size, or even just $\mathsf{poly}(N \log W) \cdot \lambda$ if computational security is enough (λ is the security parameter). They did that by describing monotone circuits that compute every threshold function, and using known monotone-circuits-to-secret-sharing compilers [BL88, VNS+03].[1] Works such as [FP12] and [Tas07] propose an explicit scheme for hierarchical threshold structures, those are solving a different (albeit somewhat related) problem than ours.

Another notable prior work is due to Zou et al. [ZMB+11], they use the Chinese Remainder Theorem to improve some efficiency parameters of weighted multi-secret sharing, but secret sizes are still the same as in the simple scheme based on Shamir sharing.

Finally, the ideas underlying our blueprint were also used in some prior works in the context of secure computation, e.g., [KMPS14, KMS16].

Organization

We present some background in section 2, then define (α, β)-ramp WSS and describe a simple rounding-based protocol for realizing it in section 3. We formulate our blueprint for WSS schemes from wiretap schemes in section 4, then describe instantiations of this blueprint from binary symmetric channels in section 5 and from additive white Gaussian noise channels in section 6.

2 Background

Notations. For an integer n, we denote $[n] = \{1, 2, \ldots, n\}$. For two distributions $\mathcal{D}, \mathcal{E}$, we denote by $SD(\mathcal{D}, \mathcal{E})$ their statistical distance. Namely $SD(\mathcal{D}, \mathcal{E}) = \frac{1}{2} \sum_{x \in X} |\mathcal{D}(x) - \mathcal{E}(x)|$, where X is the union of their support.

For a real number x and an integer η, we denote by $\lfloor x \rfloor_{2^{-\eta}}$, $\lceil x \rceil_{2^{-\eta}}$, $\lceil x \rfloor_{2^{-\eta}}$ the rounding of x down, up, or to the nearest number with precision $2^{-\eta}$, respectively. Namely, $\lfloor x \rfloor_{2^{-\eta}}$ is the largest number of the form $i/2^{\eta}$ (with i an integer) which is not larger than x, and similarly $\lceil x \rceil_{2^{-\eta}}$ is the smallest number of this form which is not smaller than x, and $\lceil x \rfloor_{2^{-\eta}}$ is one of the above which is closer to x (breaking ties arbitrarily). Omitting the $2^{-\eta}$ parameter means rounding to an integer (same as using 2^0).

2.1 Channels and Error Correcting Codes

A communication channel with input set $\mathcal{X}$ and output set $\mathcal{Y}$ is a transform that maps each input symbol $x \in \mathcal{X}$ to a distribution over the output symbols $\mathcal{Y}$. In this work we deal with additive-noise channels where $\mathcal{X} = \mathcal{Y}$ is an additive group, and the channel just adds to its input some random noise, chosen from a known distribution $\mathcal{D}$. Namely, $\mathsf{Ch} : x \mapsto x + \mathcal{D}$. We assume a memoryless channel: when sending a sequence of symbols, each symbol is transformed according to the channel Ch independently of the others (and their order is maintained).

An error-correction scheme is meant to facilitate reliable transmission of a sequence of symbols $m \in \mathcal{X}^k$ (for some k) over the channel Ch. For any input length k it consists of encoding $\mathsf{Enc} : \mathcal{X}^k \to$

[1] Those compilers essentially construct a garbled circuit for the threshold function, with the secret being the output label. Hence, they require a very large public information, namely the garbled circuit itself.

$\mathcal{X}^n$ that adds redundancy, mapping the information sequence m to a longer code-word $w \in \mathcal{X}^n$ that will be sent over the channel, and a matching decoding routine $\mathsf{Dec} : \mathcal{X}^n \to \mathcal{X}^k$ that attempts to recover the original information from the received sequence $\mathsf{Ch}(w)$. An error-correction scheme is a sequence of codes for increasing k.

The rate of a code is k/n, and the channel capacity is the highest possible rate (asymptotically as $k \to \infty$) of any scheme that achieves vanishing decoding error probability. For additive noise channels with noise distribution $\mathcal{D}$, the channel capacity is $1-h(\mathcal{D})$ where h is the Shannon entropy function. In particular, for any channel Ch and any $\nu > 0$, there exist schemes with rate ν away from capacity, in which the decoding error probability (for large enough n) is bounded below $2^{-n\cdot\mathsf{poly}(\nu)}$.

In this work we will be concerned with *Binary Symmetric Channels* (BSC, see section 5) and *Additive White Gaussian Noise* channels (AWGN, see section 6). For those channels, there exist schemes with efficient encoding/decoding procedures that approach capacity and achieve vanishing error probability. (The dependence on the slackness parameter ν = capacity-minus-rate, affects the parameters that our blueprint can achieve, and will be discussed in the sequel.)

The "more noisy" relation.

We say that a channel Ch' is more noisy than another channel Ch (or Ch is less noisy than Ch'), and denote $\mathsf{Ch} \preceq \mathsf{Ch}'$ or $\mathsf{Ch}' \succeq \mathsf{Ch}$, if there is some transform T such that $\mathsf{Ch}' = T(\mathsf{Ch})$. An example is when Ch' is obtained from Ch by adding more noise, $\mathsf{Ch}'(x) = \mathsf{Ch}(x) + \mathcal{D}$ for some noise distribution $\mathcal{D}$. It is easy to see that the capacity of Ch is at least as high as that of Ch'. Moreover, any error-correction scheme for Ch' also works for Ch.[2]

2.2 Wiretap Channel Transmission Schemes

A wiretap channel is a pair of communication channels $(\mathcal{R}, \mathcal{A})$ with the same input and output sets $\mathcal{X}, \mathcal{Y}$, where $\mathcal{R}$ is a channel from the sender to an intended receiver and $\mathcal{A}$ is the wiretap that goes to the adversary. Given a message m that the sender wants to send to the receiver, the goal is to encode it as $w = \mathsf{Enc}(m)$, so that m can be recovered (whp) from $\mathcal{R}(w)$, but not from $\mathcal{A}(w)$.

Bellare et al. defined in [BTV12] the notion of semantically secure encryption scheme for a wiretap channel (that we prefer to call a *transmission scheme*[3]). The following is essentially their definition of distinguishing security. In our setting, it is sufficient to work with what they call a "seeded" scheme, where encoding and decoding depend on a public random seed.

Definition 1 (Secure Wiretap Transmission Schemes). *Let $(\mathcal{R}, \mathcal{A})$ be a wiretap channel (for message space $\mathcal{M}$), a secure transmission scheme for it consists of (seed-dependent[4]) encoding and decoding procedures $\mathsf{Enc}_{\mathsf{sd}}, \mathsf{Dec}_{\mathsf{sd}}$ such that*

Correctness. *For all $m \in \mathcal{M}$,* $\Pr[\mathsf{Dec}_{\mathsf{sd}}(\mathcal{R}(\mathsf{Enc}_{\mathsf{sd}}(m))) = m] \geq 1 - \mathsf{negl}(|\mathsf{sd}|)$,

Secrecy. *For all $m, m' \in \mathcal{M}$, $SD\left((\mathsf{sd}, \mathcal{A}(\mathsf{Enc}_{\mathsf{sd}}(m))), (\mathsf{sd}, \mathcal{A}(\mathsf{Enc}_{\mathsf{sd}}(m')))\right) \leq \mathsf{negl}(|\mathsf{sd}|)$,*

where the probability is over the channel randomness as well as the random selection of the seed sd*, and* negl *is some negligible function.*

[2] In theory, to use a decoder for Ch' we may need to apply T to the output of $\mathsf{Ch}(w)$ before we can decode it. In practice, decoders for the high-noise Ch' always work as-is also for the low-noise Ch.

[3] This is a keyless scheme, so it differs from cryptographic encryption.

[4] We use the seed length as the security parameter for this definition.

The literature contains many constructions of wiretap channel schemes from error-correcting schemes, some of which we will be using in sections 5 and 6. For the abstract blueprint that we present in section 4, we need the "obvious" property of all the schemes in the literature, where if they work for one wiretap channel then they also work for all "easier channels." Namely, they are monotone in terms of the more-noisy relation:

Definition 2 (Monotone Schemes)**.** *A secure transmission scheme* $(\mathsf{Enc}, \mathsf{Dec})$ *for a channel* $(\mathcal{R}, \mathcal{A})$ *is* noise-monotone *if it is also a secure transmission scheme for any channel* $(\mathcal{R}', \mathcal{A}')$ *such that* $\mathcal{R}' \preceq \mathcal{R}$ *and* $\mathcal{A} \preceq \mathcal{A}'$*.*

Clearly, the secrecy condition of a transmission scheme is always monotone. The correctness condition is monotone as long as the decoding error of the underlying code is not increased by *reducing* the noise level of the channel (which is true for all coding schemes that we know of).

3 Weighted Secret Sharing

A secret-sharing scheme is a two-phase multi-party protocol for $N+1$ parties, a dealer and N shareholders. In the dealing phase, the dealer has a secret input $\boldsymbol{s}$, and it outputs a share for each shareholder, and optionally also a public share. In the reconstruction phase, a subset of the shareholders collect all their shares (and the public share if any) and attempt to use them in order to reconstruct the secret.

Each secret-sharing schemes comes with an *access structure*, consisting of a collection of qualified subsets $\Gamma \subset 2^{[N]}$ that should be able to reconstruct the secret, and a collection of *unqualified* subsets $\Psi \subset 2^{[N]}$ that should not be able to learn anything about the secret.[5] Non-perfect realizations of secret sharing come with a security parameter λ that is given as input to all the parties, and we require that the imperfections are negligible in λ.

Below we denote by $\mathsf{View}_S(\boldsymbol{s})$ the view of a subset of the shareholders $S \subset [N]$ when the secret $\boldsymbol{s}$ is shared, consisting of their own shares and the public share (if any). For a qualified set S we also denote by $\mathsf{Recover}(\mathsf{View}_S(\boldsymbol{s}))$ the value that these shareholders compute when trying to recover the secret.

Definition 3 (Secret Sharing)**.** *A secret-sharing scheme for the access structure* (Γ, Ψ) *and the space of secrets* $\mathcal{S}$*, satisfies the following (for some negligible function* $\mathsf{negl}(\cdot)$*):*

Correctness. *For any qualified subset* $S \in \Gamma$ *and any secret* $\boldsymbol{s} \in \mathcal{S}$*,*

$$\Pr[\mathsf{Recover}(\mathsf{View}_S(\boldsymbol{s})) = \boldsymbol{s}] \geq 1 - \mathsf{negl}(\lambda).$$

Secrecy. *For any unqualified subset* $T \in \Psi$ *and any two secrets* $\boldsymbol{s}, \boldsymbol{s}' \in \mathcal{S}$*, the views of* T *when sharing* $\boldsymbol{s}, \boldsymbol{s}'$ *are statistically close,*

$$SD\left(\mathsf{View}_T(\boldsymbol{s}), \mathsf{View}_T(\boldsymbol{s}')\right) \leq \mathsf{negl}(\lambda).$$

In this work we study a relaxation of threshold weighted secret sharing, (α, β)-ramp weighted secret sharing.

[5] Sometimes, but not always, we have $\Psi = \overline{\Gamma}$.

Parameters: Weights $w_1, w_2, \dots, w_N \in \mathbb{N}$, thresholds $0 < \alpha < \beta < 1$. Let $W := \sum_{i \in [N]} w_i$.

Sharing a secret $\boldsymbol{s} \in \{0,1\}^k$: Let $r = \lceil (\beta - \alpha)W \rceil$ and $k' = \lceil k/r \rceil$.

1. Break the secret into r chunks of length $\leq k'$, let $\vec{s} \in (\{0,1\}^{k'})^r$ be the resulting vector;
2. Share $\vec{s}$ using $(\alpha W, \beta W; r, W)$ multi-secret sharing, party, $j \in [N]$, gets w_j shares.

Reconstructing the secret by a qualified set S:

3. Use multi-secret reconstruction with all revealed shares to recover $\vec{s}$;
4. Concatenate the entries of $\vec{s}$ to get $\boldsymbol{s}$.

Figure 1: A rate-efficient (α, β)-ramp WSS from multi-secret sharing.

Definition 4 ((α, β)-ramp weighted secret sharing)**.** *A (α, β)-ramp weighted secret sharing for $0 < \alpha < \beta < 1$, N shareholders, and weights $w_1, \dots, w_N \in \mathbb{N}$, is a secret-sharing scheme for the access structure*

$$\Gamma = \{S \subseteq [N] : \sum_{i \in S} w_i \geq \beta W\} \text{ and } \Psi = \{T \subseteq [N] : \sum_{i \in T} w_i < \alpha W\},$$

where $W = \sum_{i \in [N]} w_i$.

Below we often use the notation $\epsilon = \beta - \alpha$ when discussing the parameters of ramp WSS schemes.

3.1 Ramp WSS from Multi-Secret Sharing

A $(T_1, T_2; r, N)$ multi-secret sharing scheme shares r secrets (from some domain) among N shareholders, with secrecy when T_1 or less of the shares are revealed and recovery when T_2 or more shares are revealed. A packed Shamir sharing, where multiple secrets are encoded in different evaluation points of a degree-$(T-1)$ polynomial, yields a $(T-r, T; r, N)$ multi-secret sharing scheme over any field of size $\geq N + r$, where each share is only a single field element. Hence it achieves a "rate" of $|\text{secret}|/|\text{share}| = r$.

This can be converted to a ramp WSS scheme using the obvious approach of giving w shares to a weight-w shareholder. This construction is described in fig. 1. To get an (α, β)-ramp WSS we need a multi-secret scheme with $N := W$, $T_1 := \alpha W$, and $T_2 := \beta W$. Using the above construction, we can pack $r = T_2 - T_1 = \epsilon W$ field elements while each underlying share is a single element.

Each shareholder in the resulting WSS scheme holds at most W shares of the underlying scheme, so we get a WSS scheme with share size $\leq W$ element that can handle secrets of size upto ϵW elements. This yields encoding rate of

$$|\text{secret}|/|\text{share}| = \frac{\epsilon W}{W} = \epsilon,$$

as long as the secret is long enough (i.e., at least ϵW field elements). This scheme is not very useful for short secrets, however, as its efficiency depends on breaking the secret into many chunks. In particular, the size of shares is still W (or more) in the worst case, regardless of how small is the secret. [6]

[6] In other contexts it is sometimes helpful to use algebraic-geometric codes instead of the Reed-Solomon codes of

Parameters: Weights $w_1, w_2, \ldots, w_N \in \mathbb{N}$, thresholds $0 < \alpha < \beta < 1$. Let $W := \sum_{i \in [N]} w_i$.

1. Let $\eta := \left\lceil \log \frac{5N}{\beta - \alpha} \right\rceil$. For all $j \in [N]$, set $w'_j := 2^\eta \cdot \left\lceil \frac{w_j}{W} \right\rceil_{2^{-\eta}}$.
2. Use the Ramp WSS from fig. 1 with the w'_j's and thresholds $\alpha' = \alpha + \frac{\beta-\alpha}{5}$ and $\beta' = \beta - \frac{\beta-\alpha}{4}$.

Figure 2: A rounded (α, β)-ramp weighted secret sharing

3.2 A Rounding-Based (α, β)-ramp WSS Protocol

We note that simple rounding can be used to roughly replace the dependence on W in the above scheme by dependent on N/ϵ. Specifically, we use the construction from fig. 1 to implement a modified version of the system, with weights that are rounded to precision of only about $(\beta - \alpha)/N$. Due to rounding errors, the modified version has a smaller gap $\epsilon' < \beta - \alpha$, but the increase can be controlled by setting the precision appropriately. Specifically, with precision of $(\beta - \alpha)/5N$ we can get $\epsilon' \geq \epsilon/2$. This simple protocol is described in fig. 2.

Lemma 1. *The protocol outline in fig. 2 is an (α, β)-ramp weighted secret sharing scheme.*

Proof. By our choice of η we get $N/2^\eta \leq (\beta - \alpha)/5$, and for every set $J \subseteq [N]$ we have

$$2^\eta \sum_{j \in J} w_j / W \;\leq\; \sum_{j \in J} w'_j \;<\; |J| + 2^\eta \sum_{j \in J} w_j / W.$$

In particular for $J = [N]$ we have $W' = \sum_{j \in [N]} w'_j \in [2^\eta, 2^\eta + N]$. For any non-qualified set $J \subseteq [N]$ with $\sum_{j \in J} w_j \leq \alpha W$ we therefore have

$$\sum_{j \in J} w'_j / W' \leq \frac{N + 2^\eta \sum_{j \in J} w_j / W}{2^\eta} \leq \frac{N + 2^\eta \cdot \alpha}{2^\eta} \leq (\beta - \alpha)/5 + \alpha = \alpha'.$$

Similarly, for any qualified set $J \subseteq [N]$ with $\sum_{j \in J} w_j \geq \beta W$ we have

$$\sum_{j \in J} w'_j / W' \geq \frac{2^\eta \sum_{j \in J} w_j / W}{N + 2^\eta} \geq \frac{\beta}{1 + (N/2^\eta)} \geq \frac{\beta}{1 + (\beta - \alpha)/5} \overset{(*)}{\geq} \beta - (\beta - \alpha)/4 = \beta'.$$

To see why inequality $(*)$ holds, note that

$$\frac{\beta}{1 + (\beta - \alpha)/5} = \frac{\beta(1 + (\beta - \alpha)/5)}{1 + (\beta - \alpha)/5} - \frac{\beta(\beta - \alpha)/5}{1 + (\beta - \alpha)/5} = \beta - \frac{\beta(\beta - \alpha)}{5 - (\beta - \alpha)} \geq \beta - \frac{\beta - \alpha}{4}.$$

□

In terms of performance for the protocol of fig. 2, the number of shares a party can receive is upper-bounded by $W' < N + 2^\eta \leq N\big(1 + \frac{10}{\beta - \alpha}\big)$. Hence, the size of shares in this scheme grows with $O(N/\epsilon)$ instead of the total weight W.

Shamir sharing, as it enables the use of smaller fields. In our case this does not seem to help, since the inefficiency comes from the number of field elements and not their size.

Sharing a secret $\boldsymbol{s} \in \{0,1\}^k$, with security parameter λ:

1. If the wiretap scheme is seeded, choose a random seed sd of length λ;
2. $\forall j \in [N]$, draw $\boldsymbol{e}_j \leftarrow \mathcal{D}_{w_j}$ and send to party j;
3. Publish sd and $\boldsymbol{g} = Enc_{\mathsf{sd}}(\boldsymbol{s}) + \sum_{j\in[N]} \boldsymbol{e}_j$.

Reconstructing the secret by a qualified set S:

Set $\boldsymbol{g}' = \boldsymbol{g} - \sum_{j\in S} \boldsymbol{e}_j$ and output $\mathsf{Dec}_{\mathsf{sd}}(\boldsymbol{g}')$.

Figure 3: The generic framework for ramp weighted secret sharing from wiretap channels.

4 A Blueprint for WSS from Wiretap Channels

Let $w_1, \ldots, w_N$ be the concrete weights that we want to implement and $0 < \alpha < \beta < 1$ be the parameters that we want to achieve. Denote $W = \sum_{i\in[N]} w_i$. An instance of our blueprint operates in some additive group $\mathcal{X}$, and consists of two components:

- A mapping from weights $w \in \mathbb{N}$ to noise distributions $\mathcal{D}_w$ over $\mathcal{X}$.
- A (seeded) noise-monotone secure transmission scheme $(\mathsf{Enc}, \mathsf{Dec})$ for a wiretap channel $(\mathcal{R}, \mathcal{A})$, such that:
  - For any qualified subset $S \subseteq [N]$ with $\sum_{i\in S} w_i \geq \beta W$, the channel $\mathcal{R}$ is more noisy than adding all the noise distributions *outside* S. Namely, $C_S \preceq \mathcal{R}$ where $C_S : x \mapsto x + \sum_{i\notin S} \mathcal{D}_{w_i}$.
  - For any unqualified subset $T \subseteq [N]$ with $\sum_{i\in S} w_i \leq \alpha W$, the channel $\mathcal{A}$ is less noisy than adding all the noise distributions *outside* T. Namely, $C_T \succeq \mathcal{A}$ where $C_T : x \mapsto x + \sum_{i\notin T} \mathcal{D}_{w_i}$.

Given these components, our WSS scheme is described in fig. 3.

Lemma 2. *If* $(\mathsf{Enc}, \mathsf{Dec})$ *is a noise-monotone secure transmission scheme for a wiretap channel* $(\mathcal{R}, \mathcal{A})$ *that satisfy the conditions above, then the scheme from fig. 3 is a secure* (α, β)*-ramp weighted secret-sharing scheme.*

Proof. This holds more or less by definition. Consider an arbitrary qualified set S and an arbitrary unqualified set T. Then by construction we have $C_S \preceq \mathcal{R}$ and $\mathcal{A} \preceq C_T$, and since $(\mathsf{Enc}, \mathsf{Dec})$ is noise-monotone then it is also a secure transmission scheme for the wiretap channel (C_S, C_T). This means on one hand that for the qualified set S, seeing $w = C_S(\mathsf{Enc}(\boldsymbol{s}))$, we have $\mathsf{Dec}(w) = \boldsymbol{s}$ with all but negligible probability. On the other hand, the unqualified set T, seeing only $C_T(\mathsf{Enc}(\boldsymbol{s}))$, cannot distinguish it from $C_T(\mathsf{Enc}(\boldsymbol{s}'))$ except with a negligible advantage. □

The public share. Our solutions, as well as some solutions from the literature (such as [BW06]), use a public share, which is known to everyone, in addition to the individual shares of the shareholders. While it is possible to eliminate the public share by adding it to each individual share, this could significantly increase the share size.[7] Instead, we chose to account for the public share separately and only count it once (rather than once per shareholder).

5 Constructions from Binary Symmetric Channels

5.1 Background

5.1.1 Binary Symmetric Channels

A binary symmetric channel (BSC) is used for sending bits. It is associated with a "crossover probability" $p \leq 1/2$, which is the probability that the received bit differs from the one that was sent. Namely, we have a Bernoulli error variable B_p with $\Pr[B_p = 1] = p$ and $\Pr[B_p = 0] = 1 - p$, and the channel is defined on message space $\{0, 1\}$ as $BSC_p : x \mapsto x + B_p \bmod 2$.

The capacity of BSC_p is $h(p)$, where h is the binary entropy function. If we are not concerned with efficient decoding, then random linear codes (with ML/MAP decoding) have rates that approach the channel capacity with exponentially small error probability. Capacity-approaching constructions with efficient decoding are known using concatenated codes [GR08] or polar codes [Ari09, GX15], with somewhat weaker bounds on the decoding error. For example, [AT09, HMTU13] show that the error probability for block-length n and rate $1 - h(p) - \nu$ is at most $\exp(-\Theta(\sqrt{n}))$, where the constant in the exponent depends on p and ν. Later results feature stronger bounds in terms of the block-length n with polynomial dependence on the slackness ν. In particular, we have

Lemma 3. *(Corollary of [BGS18, Thm 17]) For any $p < 1/2$, $\nu < 1 - h(p)$, and $\mu < 1$, there exists a code for BSC_p with rate $1 - h(p) - \nu$, block length $n = \mathsf{poly}_\mu(1/\nu)$ (for some polynomial that depends on μ), error probability $\exp(-n^\mu)$, and decoding complexity $O(n \log n)$.*

It is known that the polynomial dependence on $1/\nu$ is at least quadratic for BSC, while for some efficient constructions there is evidence that $\text{poly}_{1/2}(x) \leq x^{4.7}$ [MHU16].

5.1.2 Wiretap Schemes for Binary Symmetric Channels

Bellare et al. also described in [BTV12] a construction called ItE (Invert-then-Encode) for discrete wiretap channels, building on error-correction. The construction realizes definition 1 for the channels $(\mathcal{R}, \mathcal{A})$, using a code with low decoding error probability for $\mathcal{R}$, at a rate noticeably larger than the capacity of $\mathcal{A}$. (In particular, if the code rate approaches the capacity of $\mathcal{R}$ then this construction approaches the secrecy capacity of the wiretap channel.)

The ItE construction has integer parameters $b < k < n$ (with values as set later in this section). Identifying $\{0, 1\}^k$ with the finite field $\mathbb{F}_{2^k}$, this is a seeded construction with seed space the multiplicative group $\mathbb{F}_{2^k} \setminus \{0^k\}$ and message space $\{0, 1\}^b$. In addition, it uses error-correction encoding $\mathsf{Enc} : \{0, 1\}^k \to \{0, 1\}^n$ and the corresponding decoding $\mathsf{Dec} : \{0, 1\}^n \to \{0, 1\}^k$. The encoding and decoding routines of the ItE construction (denoted $\mathsf{Enc}'_{\mathsf{sd}}, \mathsf{Dec}'_{\mathsf{sd}}$) are described in fig. 4. The following is a re-phrasing of Lemma 5.3 and Lemmas 5.5-5.6 from [BT12]:

Lemma 4. *([BT12, Lemma 5.3]) If $(\mathsf{Enc}, \mathsf{Dec})$ is an error-correction scheme with decoding-error probability at most ϵ for the channel $\mathcal{R}$, then the ItE scheme $(\mathsf{Enc}_{\mathsf{sd}}, \mathsf{Dec}_{\mathsf{sd}})$ from fig. 4 is correct for $(\mathcal{R}, \mathcal{A})$ with correction holding with probability $\geq 1 - \epsilon$.* □

[7] In our solutions it will only double the share size, but in some solutions in the literature the public share is much larger than the individual shares.

```
Encoding: $\mathsf{Enc}'_{\mathsf{sd}}(M \in \{0,1\}^b)$ with seed $\mathsf{sd} \in \mathbb{F}_{2^k} \setminus \{0^k\}$:
    1. Choose a random $R \leftarrow \{0,1\}^{k-b}$, let $Y = (M|R) \in \mathbb{F}_{2^k}$ be the concatenation;
    2. Set $X := Y/\mathsf{sd} \in \mathbb{F}_{2^k}$;
    3. Send the message $W := \mathsf{Enc}(X) \in \{0,1\}^n$.
Decoding: $\mathsf{Dec}'_{\mathsf{sd}}(W' \in \{0,1\}^n)$:
    1. Use error-correction to get $X' = \mathsf{Dec}(W')$;
    2. Compute $Y' := X' \cdot \mathsf{sd}$ ;
    3. Output $M'$, the first $b$ bits of $Y'$.
```

Figure 4: The ItE construction from [BTV12].

Lemma 5. *(Corollary of [BT12, Lemmas 5.5-5.6]) Let $\mathcal{A}$ be a symmetric memoryless channel with capacity $c(\mathcal{A})$. Assume that $\frac{k}{n}$ (the rate of Enc) is larger than $c(\mathcal{A})$, denote the slackness by $\rho = \frac{k}{n} - c(\mathcal{A})$, and let λ be the security parameter. Then for any $0 < \delta < \rho - \frac{2\lambda}{n}$, setting $b := \lfloor n(\rho - \delta) - 2\lambda \rfloor$ in the ItE construction yields a wiretap transmission scheme with secrecy upto statistical distance $4 \cdot 2^{-\delta^2 n/11} + 2 \cdot 2^{-\lambda}$.*

Plugging the coding parameter from above, we get the following instantiation:

Corollary 6. *For a binary symmetric wiretap channel (BSC_{p_R}, BSC_{p_A}) with $0 \le p_R < p_A < 1/2$, denote $\xi := h(p_A) - h(p_R)$. There exists an instance of the ItE scheme $(\mathsf{Enc}_{\mathsf{sd}}, \mathsf{Dec}_{\mathsf{sd}})$ with security parameter λ and*

- *Encoding size $n = \max\left(\mathsf{poly}_{\frac{1}{2}}(\frac{4}{\xi}), \lambda^2, \frac{44\lambda}{\xi^2}\right)$;* [8]
- *Seed space $\mathbb{F}_{2^k} \setminus \{0^k\}$ with $k = (1 - h(p_A) + \frac{3\xi}{4})n = (1 - h(p_R) - \frac{\xi}{4})n$; and*
- *Message space $\{0,1\}^b$, $b \ge (\frac{\xi}{4} - \frac{2}{\lambda})n$;*

such that

- *For all $m \in \{0,1\}^b$, $\Pr[\mathsf{Dec}_{\mathsf{sd}}(BSC_{p_R}(\mathsf{Enc}_{\mathsf{sd}}(m))) = m] \ge 1 - 2^{-\lambda}$;*
- *For all $m, m' \in \{0,1\}^b$,*

$$SD\big((\mathsf{sd}, BSC_{p_A}(\mathsf{Enc}_{\mathsf{sd}}(m))),\ (\mathsf{sd}, BSC_{p_A}(\mathsf{Enc}_{\mathsf{sd}}(m')))\big) \le 6 \cdot 2^{-\lambda}.$$

Proof. Recall that k determines both the seed space of the ItE construction and the input space for the underlying error-correcting code. The rate of the underlying code is therefore $k/n = 1 - h(p_R) - \frac{\xi}{4}$, and by lemma 3 we can find such codes as soon as the encoding-length exceeds

[8] $\mathsf{poly}_{\frac{1}{2}}$ is the polynomial from lemma 3 for $\mu = \frac{1}{2}$.

Parameters:
Weights $w_1, w_2, \ldots, w_N \in \mathbb{N}$, thresholds $0 < \alpha < \beta < 1$, security parameter λ.

- Let $W := \sum_{i\in[N]} w_i$ and $\gamma := \frac{\beta-\alpha}{2(1-\alpha)^2}$. We assume that $\gamma \le \frac{1}{2}$. [a]
- Let $p_R := \gamma(1-\beta)$ and $p_A := \gamma(1-\alpha) - \gamma^2(1-\alpha)^2$.
  (In section 5.3 below we show that $\frac{1}{4} \ge p_A \ge p_R + (\beta-\alpha)^2/4$.)
- Let $(\mathsf{Enc}, \mathsf{Dec})$ (with parameters n, k, b) be as in the ItE construction from corollary 6 for the wiretap channel (BSC_{p_R}, BSC_{p_A}).

Sharing a secret $\boldsymbol{s} \in \{0,1\}^k$:

1. $\forall j \in [N]$, set $p_j := \gamma \cdot \frac{w_j}{W}$, draw $\boldsymbol{e}_j \leftarrow (\mathsf{Bernoulli}_{p_i})^n$ and send to party j;
2. Draw a random $\mathsf{sd} \in \mathbb{F}_{2^k} \setminus \{0^k\}$, publish sd and $\boldsymbol{g} := \mathsf{Enc}_{\mathsf{sd}}(\boldsymbol{s}) + \sum_{j\in[N]} \boldsymbol{e}_j \bmod 2$.

Reconstructing the secret by a qualified set S:
Set $\boldsymbol{g}' = \boldsymbol{g} + \sum_{j\in S} \boldsymbol{e}_j \bmod 2$ and output $\mathsf{Dec}_{\mathsf{sd}}(\boldsymbol{g}')$.

[a] This holds if α, β are not too big. For example if $\alpha \le \frac{1}{2}$ and $\beta \le \frac{1+\alpha}{2}$.

Figure 5: Weighted secret sharing from symmetric binary wiretap channels.

$\mathsf{poly}_{1/2}(4/\xi)$, with decoding error probability at most $\exp(-n^{1/2}) < 2^{-\sqrt{n}}$. If $n \ge \lambda^2$ then this is bounded below $2^{-\lambda}$, and due to lemma 4 the same holds for correctness of the ItE construction.

For the secrecy part, we have rate $k/n = 1 - h(p_A) + \frac{3\xi}{4}$, and we use $\delta = \frac{\xi}{2}$ in lemma 5. This yields $b = \frac{\xi}{4}n - 2\lambda$, and since $n \ge \lambda^2$ then $b \ge (\frac{\xi}{4} - \frac{2}{\lambda})n$. If we also have $n \ge \frac{44\lambda}{\xi^2}$, then $\delta^2 n/11 \ge (\xi/2)^2 \cdot (44\lambda/\xi^2)/11 = \lambda$, and therefore the statistical distance is bounded by

$$4 \cdot 2^{-\delta^2 n/11} + 2 \cdot 2^{-\lambda} \le 4 \cdot 2^{-\lambda} + 2 \cdot 2^{-\lambda} = 6 \cdot 2^{-\lambda}.$$

□

Remark. Different from most works in the literature, in the setting above we do not aim at achieving the secrecy capacity in the limit. Rather, we try to maintain a small encoding size n relative not just the message size b, but also the security parameter λ and the parameters p_R, p_A. [9]

5.2 Our Construction

In fig. 5 we show how to use the ItE instance from corollary 6 to get an (α, β)-ramp WSS for given weights $w_1, w_2, \ldots, w_N$ and thresholds $0 < \alpha < \beta$. (The parameters below are chosen for α, β that are not very close to one, but they can be modified to handle larger α, β at a small performance loss.)

Clearly, this construction is an instance of the blueprint from fig. 3, instantiated over the additive group $\mathbb{F}_{2^k}$, using the noise distributions $D_w = \mathsf{Bernoulli}_{\gamma w/W}$ and the ItE construction from

[9] In particular, we opted for losing a constant factor in the ratio b/n in return for better dependency on λ and ξ.

corollary 6 for the wiretap channel (CSB_{p_R}, CSB_{p_A}). It is also clear that the ItE construction is noise-monotone (since the underlying error-correction codes are).

The only thing left to prove in order to use lemma 2, is that for any qualified S and unqualified T, the corresponding channels satisfy $C_S \preceq BSC_{p_R}$ and $C_T \succeq BSC_{p_A}$. To that end, we use the following technical lemma (whose proof is in the appendix)

Lemma 7. *Let $\mathcal{B}_1, \mathcal{B}_2, ..., \mathcal{B}_t$ be independent Bernoulli random variables, and denote $\mathcal{S} := \sum_{j\in[t]} \mathcal{B}_j$ and $\epsilon := \sum_{j\in[t]} \Pr[\mathcal{B}_j = 1]$. Then $\Pr[\mathcal{S} \text{ is odd}] \in [\epsilon - \epsilon^2,\ \epsilon]$.*

We can now complete the proof that the ItE-based construction above satisfies all the conditions of lemma 2.

Lemma 8. *With the parameters as set in fig. 5:*

(A) For every subset $S \subseteq [N]$ with $\sum_{j\in S} w_j \geq \beta W$, we have $C_S \preceq BSC_{p_R}$ where $C_S : x \mapsto x + \sum_{j\notin S} \mathsf{Bernoulli}_{p_j}$.

(B) For every subset $T \subseteq [N]$ with $\sum_{j\in S} w_j \leq \alpha W$, we have $C_T \succeq BSC_{p_A}$ where $C_T : x \mapsto x + \sum_{j\notin T} \mathsf{Bernoulli}_{p_j}$.

Proof. Clearly C_S, C_T are memoryless binary symmetric channels, so all we need to show is that the crossover probability of C_S is at most p_R and the that of C_T is at least p_A (and they are less than $\frac{1}{2}$). For any bit position $\ell \in [n]$, the $\boldsymbol{e}_j[\ell]$'s are independent Bernoulli random variables, with success probabilities $p_j := \Pr[\boldsymbol{e}_j[\ell] = 1] = \gamma w_j / W$.

For the channel C_S, denote $P_{\overline{S}} := \sum_{j\notin S} p_j = \gamma(\sum_{j\notin S} w_j)/W$. The crossover probability of C_S is exactly the probability that an odd number of these variables $\boldsymbol{e}_j[\ell]$ are set to one, which by lemma 7 is at most $P_{\overline{S}}$. Since S is a qualified set then we know that $P_{\overline{S}} = \gamma(\sum_{j\notin S} w_j)/W \leq \gamma(1-\beta) = p_R < \frac{1}{2}$, as needed. This implies that $C_S \preceq BSC_{p_R}$.

Similarly, for the channel C_T, denote $P_{\overline{T}} := \sum_{j\notin T} p_j = \gamma(\sum_{j\notin T} w_j)/W$. The crossover probability of C_T is exactly the probability that an odd number of these variables $\boldsymbol{e}_j[\ell]$ are set to one, which by lemma 7 is at least $P_{\overline{T}} - P_{\overline{T}}^2$. Since T is an unqualified set then we know that $P_{\overline{T}} = \gamma(\sum_{j\notin T} w_j)/W \geq \gamma(1-\alpha)$. Moreover, since we assume that $\gamma \leq \frac{1}{2}$ then also $P_{\overline{T}} \leq \gamma \leq \frac{1}{2}$, and since $f(x) = x - x^2$ is monotonically increasing in the range $(0, \frac{1}{2})$, then from $P_{\overline{T}} \leq \gamma(1-\alpha)$ we conclude that $P_{\overline{T}} - P_{\overline{T}}^2 \geq \gamma(1-\alpha) - \gamma^2(1-\alpha)^2 = p_A$. This implies that $C_T \succeq BSC_{p_A}$. □

5.3 Performance Characteristics of This Construction

Recall that we set the scaling parameter γ and the probabilities p_R, p_A as

$$\gamma := \frac{\beta - \alpha}{2(1-\alpha)^2}, \quad p_R = \gamma(1-\beta), \quad \text{and } p_A = \gamma(1-\alpha) - \gamma^2(1-\alpha)^2.$$

We can bound the gap $p_A - p_R$ by

$$\begin{aligned} p_A - p_R &= \gamma\big((1-\alpha) - \gamma(1-\alpha)^2 - (1-\beta)\big) = \gamma\big(\beta - \alpha - \gamma(1-\alpha)^2\big) \\ &= \frac{\beta-\alpha}{2(1-\alpha)^2} \cdot \left(\beta - \alpha - \frac{\beta-\alpha}{2(1-\alpha)^2}(1-\alpha)^2\right) = \frac{\beta-\alpha}{2(1-\alpha)^2} \cdot \frac{\beta-\alpha}{2} \\ &= \frac{(\beta-\alpha)^2}{4(1-\alpha)^2} > \frac{(\beta-\alpha)^2}{4}. \end{aligned} \tag{1}$$

Below it will be convenient to denote $\epsilon = \beta - \alpha$, and by equation (1) above we have $p_A - p_R \geq \epsilon^2/4$. Moreover, we know that $p_A \leq \frac{1}{4}$ (since $x - x^2 \leq \frac{1}{4}$ for every $x \in \mathbb{R}$). We thus have $0 < p_R < p_A \leq \frac{1}{4}$, and the binary entropy function $h(\cdot)$ has derivative greater than one in the region $(0, \frac{1}{4}]$, we get

$$\xi := h(p_A) - h(p_R) > p_A - p_R > \epsilon^2/4.$$

By corollary 6, there is a transmission scheme $(\mathsf{Enc_{sd}}, \mathsf{Dec_{sd}})$ for the wiretap channel (BSC_{p_R}, BSC_{p_A}) with correctness/secrecy upto $O(2^{-\lambda})$ and parameters

- **Encoding length:** $n \leq \max\left(\mathsf{poly}_{\frac{1}{2}}(\frac{16}{\epsilon^2}), \lambda^2, \frac{704\lambda}{\epsilon^4}\right)$;
- **Message length:** $b \geq (\frac{\epsilon^2}{16} - \frac{2}{\lambda})n$;
- **Seed length:** $k = (1 - h(p_A) + \frac{3\xi}{4})n$.

Recall that for this scheme, we have secrets of length b, each shareholder gets a share of length n, and the public share is of size $n + k$. Note also that n, k, b depend only the thresholds α, β and not on the weights themselves. Thus, we get a scheme where the share sizes are independent of the weights.

When the gap $\epsilon = \beta - \alpha$ is a constant, the share sizes are just a constant factor larger than the secret size b. As the gap gets smaller, the share sizes grow as a polynomial in $1/\epsilon$, but it is a rather large polynomial (at least quartic, with coefficients in the hundreds). To do better, we explore in the next section a different instantiation of our blueprint, using (quantized) additive Gaussian noise channels.

6 WSS from AWGN Wiretap Channels

6.1 Background

6.1.1 Additive White Gaussian Noise Channels

Additive white Gaussian noise channels (AWGN) communicate real numbers rather than bits. For each symbol $x \in \mathbb{R}$ transmitted over the channel, the received symbol is $y = x + e$ (addition over the reals), where e is a zero-mean Normal random variable. The variance σ^2 of e is the noise level of the channel.

Symbols transmitted over the channel are chosen subject to some power constraint, specifically their (expected) square is bounded by the *power parameter* P of the sender. (For example we could use $\pm\sqrt{P}$ to encode bits.) The quality of the channel is determined by the ratio between the power and the noise, called the signal-to-noise ratio: $SNR = P/\sigma^2$. [10] Below it will be convenient to fix the power to $P = 1$ and set the variance accordingly. We denote the AWGN channel with variance σ^2 (and power $P = 1$) by $AWGN_{\sigma^2} : x \mapsto x + \mathcal{N}(0, \sigma^2)$. The capacity of this channel (denoted $c(\sigma)$ below) is

$$c(\sigma) := \mathsf{capacity}(AWGN_{\sigma^2}) = \ln\left(1 + \frac{1}{\sigma^2}\right).$$

(The general formula is $\ln\left(1 + \frac{P}{\sigma^2}\right)$ but we are fixing $P = 1$.) There are known constructions of error-correcting codes with efficient decoding for the AWGN that approach capacity, see for

[10] Clearly, scaling P and σ^2 by the same factor has no effect on the channel quality.

example [EZ04, LYLW19]. AWGN codes can achieve somewhat better performance than BSC codes, since they use "soft decoding" vs. the "hard decoding" that's inherent in BSC code. To the best of our knowledge, however, this improvement has little effect on their asymptotic behavior.[11] For our purposes, we therefore only assume the following parameters, which are the same as we have for the BSC case:

Assumption 9. *For any $\sigma \in \mathbb{R}$ and slackness parameter $\nu < c(\sigma)$, there exists a code for $AWGN_{\sigma^2}$ with rate $c(\sigma) - \nu$, block length $n = \mathsf{poly}(1/\nu)$ (for some polynomial), error probability $\exp(-\sqrt{n})$, and decoding complexity polynomial in n.*

6.1.2 AWGN Wiretap Channels

Tyagi and Vardy described in [TV14] a modular construction (in the same spirit as [BTV12]) that combines AWGN codes with randomness extractors. If the underlying code approaches the receiver channel capacity, then the Tyagi-Vardy scheme can be made to approach the secrecy capacity of the wiretap channel. A different approach for a secrecy-capacity-approaching schemes was provided by Liu et al. [LYL18].

These AWGN constructions may be practically more efficient than their BSC counterparts, but as far as we know the improvement has little effect on their asymptotic behavior. We therefore assume the following (which is a counterpart of corollary 6):

Assumption 10. *For an AWGN wiretap channel $(AWGN_{\sigma_r^2}, AWGN_{\sigma_a^2})$ with $0 \le \sigma_r < \sigma_a$, denote $\xi := c(\sigma_r) - c(\sigma_a)$. There exists a seeded wiretap transmission scheme $(\mathsf{Enc_{sd}}, \mathsf{Dec_{sd}})$ with security parameter λ and*

- *Encoding size $n = \max\big(\mathsf{poly}(\frac{1}{\xi}), \lambda^2, O(\frac{\lambda}{\xi^2})\big)$;*
- *Seed size $k = (c(\sigma_a) + \Theta(\xi))n = (c(\sigma_r) - \Theta(\xi))n$; and*
- *Message space $\{0,1\}^b$, $b \ge (\Theta(\xi) - \frac{2}{\lambda})n$;*

such that

- *For all $m \in \{0,1\}^b$, $\Pr[\mathsf{Dec_{sd}}(AWGN_{\sigma_r^2}(\mathsf{Enc_{sd}}(m))) = m] \ge 1 - 2^{-\lambda}$;*
- *For all $m, m' \in \{0,1\}^b$,*

$$SD\big((\mathsf{sd}, AWGN_{\sigma_a^2}(\mathsf{Enc_{sd}}(m))),\ (\mathsf{sd}, AWGN_{\sigma_a^2}(\mathsf{Enc_{sd}}(m')))\big) \le 2^{-\lambda}.$$

6.2 Our Construction

In fig. 6 we show how to use the wiretap scheme from assumption 10 to get an (α, β)-ramp WSS for given weights $w_1, w_2, \ldots, w_N$ and thresholds $0 < \alpha < \beta < 1$. This is more or less a direct application of our blueprint, except that we need to describe how to quantize the various real numbers to finite precision.

This construction is essentially an instance of the blueprint from fig. 3, instantiated over the real numbers, using the noise distributions $D_w = \mathcal{N}(0, w_j/W)$ and the wiretap scheme from corollary 6

[11] It may mean smaller required block-length for the same gap-to-capacity, but the dependence on the gap is still polynomial.

Parameters:

Weights $w_1, w_2, \ldots, w_N \in \mathbb{N}$, thresholds $0 < \alpha < \beta < 1$, security parameter λ.

- Let $W := \sum_{i\in[N]} w_i$, $\sigma_r = \sqrt{1-\beta}$, and $\sigma_a = \sqrt{1-\alpha}$.
- Let $(\mathsf{Enc}, \mathsf{Dec})$ be as in assumption 10 for the wiretap channel $(AWGN_{\sigma_r^2}, AWGN_{\sigma_a^2})$, with parameters n, k, b.
- Let $\eta \in \mathbb{N}$ be a *quantization parameter*, to be determined below.

Sharing a secret $\boldsymbol{s} \in \{0,1\}^k$:

1. $\forall j \in [N]$, set $\sigma_j := \sqrt{w_j/W}$, draw $\boldsymbol{e}'_j \leftarrow \mathcal{N}(0, \sigma_j^2)^n$. If $\|\boldsymbol{e}'_j\|_\infty > \lambda$ then abort. Else round it to precision $2^{-\eta}$ to get $\boldsymbol{e}_j := \lceil \boldsymbol{e}'_j \rfloor_{2^{-\eta}}$ and send to party j;
2. Draw $\mathsf{sd} \leftarrow \{0,1\}^k \setminus \{0^k\}$ and set $\boldsymbol{g} := \left\lceil \mathsf{Enc}_{\mathsf{sd}}(\boldsymbol{s}) - \sum_{j\in[N]} \boldsymbol{e}'_j \right\rfloor_{2^{-\eta}}$. (Note that the summation is done on the original $\boldsymbol{e}'_j$'s, and only the result is rounded.)

   Abort if $\|\boldsymbol{g}\|_\infty > \lambda$, else publish $\boldsymbol{g}$ and the seed sd.

Reconstructing the secret by a qualified set S:

Set $\boldsymbol{g}' = \boldsymbol{g} + \sum_{j\in S} \boldsymbol{e}_j$ and output $\mathsf{Dec}_{\mathsf{sd}}(\boldsymbol{g}')$.

Figure 6: Weighted secret sharing from additive white Gaussian noise wiretap channels.

for $(AWGN_{\sigma_r^2}, AWGN_{\sigma_a^2})$. (As before it is clear the construction is noise-monotone since the underlying error-correction codes are).

The slight deviations from the blueprint, is that the dealer aborts if the errors are too big, and that the various quantities are rounded. For the first aspect, we note that $\sigma_j \le 1$ for all j (and even their sum is less than one), therefore the abort probability is exponentially small in λ. The effects of rounding are analyzed in section 6.3 below, but first we prove that otherwise we would get all the conditions that we need.

That is, we consider an idealized protocol (that cannot be implemented) where we remove the rounding steps and use all the quantities with infinite precision We prove that in that case, for any qualified S and unqualified T, the corresponding channels satisfy $C_S \preceq AWGN_{\sigma_r^2}$ and $C_T \succeq AWGN_{\sigma_a^2}$.

Lemma 11. *With the parameters as set in fig. 6, but excluding the rounding steps, we have:*

(A) For every subset $S \subseteq [N]$ with $\sum_{j\in S} w_j \ge \beta W$, we have $C_S \preceq AWGN_{\sigma_r^2}$ where $C_S : x \mapsto x + \sum_{j\notin S} \mathcal{N}(0, \sigma_j^2)$.

(B) For every subset $T \subseteq [N]$ with $\sum_{j\in S} w_j \le \alpha W$, we have $C_T \succeq AWGN_{\sigma_a^2}$ where $C_T : x \mapsto x + \sum_{j\notin T} \mathcal{N}(0, \sigma_j^2)$.

Proof. These properties follow by definition, due to the additive nature of Normal random variables. Indeed, for every set $J \subseteq [N]$ (qualified or not), denote $\sigma_J := \sqrt{\sum_{j\notin J} \sigma_j^2}$. Then $\sum_{j\notin J} \mathcal{N}(0, \sigma_j^2) = \mathcal{N}\big(0, \sigma_J^2\big)$, and the corresponding channel is $C_J : x \mapsto x + \mathcal{N}(0, \sigma_J^2)$.

For a qualified set S we have $\sum_{j\in S} w_j/W \geq \beta$ and therefore $\sigma_S^2 \leq 1-\beta = \sigma_r^2$ and $C_S \preceq AWGN_{\sigma_r^2}$. Similarly, for an unqualified set T we have $\sum_{j\in T} w_j/W \leq \alpha$ and therefore $\sigma_T^2 \geq 1-\alpha = \sigma_a^2$ and $C_T \succeq AWGN_{\sigma_a^2}$. □

6.3 The Effect of Quantization

As seen in fig. 6, a dealer must round all $\boldsymbol{e}'_j$ to precision $2^{-\eta}$. The security of the protocol outlined in fig. 6 still holds with rounding though.

Lemma 12. *If an adversary can break the protocol outlined in fig. 6, an adversary can break the Gaussian Wiretap Channel.*

Proof. If an adversary could break the security of the protocol with attack A, then an adversary could recover the secret from an unqualified set, $J \subseteq [N]$, when the dealer does not round. This can be done by the adversary trivially rounding all $\boldsymbol{e}'_j$, for all $j \in J$, and then use A to recover the secret. By lemma 11, we can see that the security of the protocol without rounding is equal to that of the Gaussian wiretap channel. □

Lemma 13. *Correctness can be unaffected when $\eta = \lambda + \log N$.*

Proof. Rounding one share $\boldsymbol{e}'_j$ may induce an error of up to $2^{-\eta}$. So, for any set $S \subseteq [N]$ there is at most $N2^{-\eta}$ error. Thus, we have at most $N \cdot \frac{1}{N} \cdot 2^{-\lambda} = 2^{-\lambda}$ error. Then, we note that the variance of noise increases by at most $2^{-2\lambda}$. Let Y be the random variable associated with the summed quantization error and X be the random variable associated with a sample of $\mathcal{N}(0, \sigma^2)$.

Then,

$$\begin{aligned}
V[X+Y] &= \mathbf{E}\left[(X+Y)^2\right] + \mathbf{E}\left[X+Y\right]^2 \\
&= \mathbf{E}\left[(X+Y)^2\right] = \sigma^2 + 2\,\mathbf{E}[XY] + \mathbf{E}[Y]^2 \\
&\leq \sigma^2 + 2\sum_{x,y} xy \Pr[X=x] + 2^{-2\lambda} \\
&\leq \sigma^2 + 2\cdot 2^{-\lambda}\sum_{x,y} x \Pr[x] + 2^{-2\lambda} = \sigma^2 + 2^{-2\lambda}.
\end{aligned}$$

Thus, if the decoder works with $2^{-2\lambda}$ higher noise variance, the decoder should still be able to decode the quantized shares. □

6.4 Performance Characteristics of This Construction

As described in fig. 6, the real numbers used in this scheme are all smaller than λ and use η bits of precision to the right of the binary point. Hence, the total size of the individual shares is $n(\eta + \log\lambda)$, and the public share is of size $k + n(\eta + \log\lambda)$. Setting $\eta = \lambda + \log N$ as in section 6.3, we therefore get share sizes $O(n(\lambda + \log N))$ for secrets of size $b = \Omega(n)$ bits, regardless of the weights.

The code-length n depends polynomially on the security parameter λ and on $1/\epsilon$ (where $\epsilon := \beta - \alpha$). For AWGN we can get $\xi = c(p_r) - c(p_a) = \Theta(\epsilon)$ (vs $\xi = \Theta((\beta-\alpha)^2)$ for BSC), since do not need the scale-down factor γ. Hence, we get better polynomial dependence on $\beta - \alpha$, even if we don't take into account the better rates of AWGN codes as compared to BSC codes with the same power. On the other hand, the presence of rounding errors means that the share sizes have some dependence on N, to the tune of a $\log N$ multiplicative factor.

6.5 Discussion: An Alternative Realization Using Discrete Gaussians

An Alternative to rounding Normal variables to finite precision would be to draw everything from discrete Gaussian distribution of various parameters. This has an advantage of eliminating rounding errors (and therefore the dependence on N). But to keep accuracy the parameters must be set so that even the smallest Gaussian parameter weight is sufficiently larger than the smoothing parameter of the underlying lattice (i.e. the precision), bringing in at least a logarithmic dependence on the actual weights w_i. This can be handled by pre-scaling all the weights to accuracy of (slightly more than) $1/\epsilon$, but that will induce rounding errors again, and a (logarithmic) dependence on the number of parties.

7 Conclusions

In this work, we study a relaxation for weighted secret sharing with a gap between the qualified and unqualified sets, and described two different types of constructions, one based on rounding and the other using a new connection to wiretap schemes. Both types have share size independent of total weight, and dependent only the gap between qualified and unqualified sets. We expect that our wiretap-based constructions can perhaps be improved, maybe using better codes. In particular there is no reason that we know of why the AWGN construction cannot achieve linear dependence on the gap, namely shares of size $O(1/(\beta - \alpha))$.

References

[Ari09] Erdal Arikan. Channel polarization: A method for constructing capacity-achieving codes for symmetric binary-input memoryless channels. *IEEE Transactions on Information Theory*, 55(7):3051–3073, 2009.

[AT09] Erdal Arikan and Emre Telatar. On the rate of channel polarization. In *2009 IEEE International Symposium on Information Theory*, pages 1493–1495, 2009.

[BGS18] Jaroslaw Blasiok, Venkatesan Guruswami, and Madhu Sudan. Polar codes with exponentially small error at finite block length. In Eric Blais, Klaus Jansen, José D. P. Rolim, and David Steurer, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2018*, volume 116 of *LIPIcs*, pages 34:1–34:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018.

[BL88] Josh Cohen Benaloh and Jerry Leichter. Generalized secret sharing and monotone functions. In Shafi Goldwasser, editor, *Advances in Cryptology - CRYPTO'88*, volume 403 of *Lecture Notes in Computer Science*, pages 27–35. Springer, 1988.

[Bla79] George R. Blakley. Safeguarding cryptographic keys. In *1979 International Workshop on Managing Requirements Knowledge (MARK)*, pages 313–318, 1979.

[BM84] G. R. Blakley and Catherine A. Meadows. Security of ramp schemes. In G. R. Blakley and David Chaum, editors, *Advances in Cryptology, Proceedings of CRYPTO '84, Santa Barbara, California, USA, August 19-22, 1984, Proceedings*, volume 196 of *Lecture Notes in Computer Science*, pages 242–268. Springer, 1984.

[BT12] Mihir Bellare and Stefano Tessaro. Polynomial-time, semantically-secure encryption achieving the secrecy capacity. *IACR Cryptology ePrint Archive*, page 22, 2012.

[BTV12] Mihir Bellare, Stefano Tessaro, and Alexander Vardy. Semantic security for the wiretap channel. In Reihaneh Safavi-Naini and Ran Canetti, editors, *Advances in Cryptology - CRYPTO 2012 - 32nd Annual Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2012. Proceedings*, volume 7417 of *Lecture Notes in Computer Science*, pages 294–311. Springer, 2012. Also available from https://arxiv.org/abs/1201.2205.

[BTW05] Amos Beimel, Tamir Tassa, and Enav Weinreb. Characterizing ideal weighted threshold secret sharing. In *Theory of Cryptography Conference*, pages 600–619. Springer, 2005.

[BW06] Amos Beimel and Enav Weinreb. Monotone circuits for monotone weighted threshold functions. *Information Processing Letters*, 97(1):12–18, 2006.

[CDN15] Ronald Cramer, Ivan Damgård, and Jesper Buus Nielsen. *Secure Multiparty Computation and Secret Sharing.* Cambridge University Press, 2015.

[EZ04] U. Erez and R. Zamir. Achieving 1/2 log(1+SNR) on the AWGN channel with lattice encoding and decoding. *IEEE Transactions on Information Theory*, 50(10):2293–2314, 2004.

[FP12] Oriol Farras and Carles Padró. Ideal hierarchical secret sharing schemes. *IEEE transactions on information theory*, 58(5):3273–3286, 2012.

[GR08] Venkatesan Guruswami and Atri Rudra. Concatenated codes can achieve list-decoding capacity. In *Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '08, page 258–267, USA, 2008. Society for Industrial and Applied Mathematics.

[GX15] Venkatesan Guruswami and Patrick Xia. Polar codes: Speed of polarization and polynomial gap to capacity. *IEEE Trans. Inf. Theory*, 61(1):3–16, 2015.

[HMTU13] S. Hamed Hassani, Ryuhei Mori, Toshiyuki Tanaka, and Rüdiger L. Urbanke. Rate-dependent analysis of the asymptotic behavior of channel polarization. *IEEE Transactions on Information Theory*, 59(4):2267–2276, 2013.

[KMPS14] Daniel Kraschewski, Hemanta K. Maji, Manoj Prabhakaran, and Amit Sahai. A full characterization of completeness for two-party randomized function evaluation. In Phong Q. Nguyen and Elisabeth Oswald, editors, *Advances in Cryptology - EUROCRYPT 2014 - 33rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Copenhagen, Denmark, May 11-15, 2014. Proceedings*, volume 8441 of *Lecture Notes in Computer Science*, pages 659–676. Springer, 2014.

[KMS16] Dakshita Khurana, Hemanta K. Maji, and Amit Sahai. Secure computation from elastic noisy channels. In Marc Fischlin and Jean-Sébastien Coron, editors, *Advances in Cryptology - EUROCRYPT 2016 - 35th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Vienna, Austria, May 8-12, 2016, Proceedings, Part II*, volume 9666 of *Lecture Notes in Computer Science*, pages 184–212. Springer, 2016.

[LYL18] Ling Liu, Yanfei Yan, and Cong Ling. Achieving secrecy capacity of the gaussian wiretap channel with polar lattices. *IEEE Transactions on Information Theory*, 64(3):1647–1665, 2018. Also available at https://arxiv.org/abs/1503.02313.

[LYLW19] Ling Liu, Yanfei Yan, Cong Ling, and Xiaofu Wu. Construction of capacity-achieving lattice codes: Polar lattices. *IEEE Transactions on Communications*, 67(2):915–928, 2019. Also available at https://arxiv.org/abs/1411.0187.

[MHU16] Marco Mondelli, S. Hamed Hassani, and Rüdiger L. Urbanke. Unified scaling of polar codes: Error exponent, scaling exponent, moderate deviations, and error floors. *IEEE Transactions on Information Theory*, 62(12):6698–6712, 2016.

[Sha79] Adi Shamir. How to share a secret. *Communications of the ACM*, 22(11):612–613, 1979.

[Tas07] Tamir Tassa. Hierarchical threshold secret sharing. *Journal of cryptology*, 20(2):237–264, 2007.

[TV14] Himanshu Tyagi and Alexander Vardy. Explicit capacity-achieving coding scheme for the gaussian wiretap channel. In *2014 IEEE International Symposium on Information Theory*, pages 956–960, 2014. See also https://arxiv.org/abs/1412.4958.

[VNS$^+$03] Vinod Vaikuntanathan, Arvind Narayanan, K. Srinathan, C. Pandu Rangan, and Kwangjo Kim. On the power of computational secret sharing. In Thomas Johansson and Subhamoy Maitra, editors, *Progress in Cryptology - INDOCRYPT 2003*, volume 2904 of *Lecture Notes in Computer Science*, pages 162–176. Springer, 2003.

[ZMB$^+$11] Xukai Zou, Fabio Maino, Elisa Bertino, Yan Sui, Kai Wang, and Feng Li. A new approach to weighted multi-secret sharing. In *2011 Proceedings of 20th International Conference on Computer Communications and Networks (ICCCN)*, pages 1–6. IEEE, 2011.

A Bound on Odd Number of Heads

Lemma 7. *Let $\mathcal{B}_1, \mathcal{B}_2, ..., \mathcal{B}_t$ be independent Bernoulli random variables, and denote $\mathcal{S} := \sum_{j\in[t]} \mathcal{B}_j$ and $\epsilon := \sum_{j\in[t]} \Pr[\mathcal{B}_j = 1]$. Then*

$$\Pr[\mathcal{S} \text{ is odd}] \in [\epsilon - \epsilon^2,\ \epsilon].$$

Proof. The upper bound follows immediately from the union bound:

$$\Pr[\mathcal{S} \text{ is odd}] \le \Pr[\mathcal{S} \ge 1] \le \sum_{j\in[t]} \Pr[\mathcal{B}_j = 1] = \epsilon.$$

Once we have the upper bound, we can prove the lower bound by induction on t. For the base case $t = 1$, we have $\mathcal{S} = \mathcal{B}_1$ and therefore $\Pr[\mathcal{S} \text{ is odd}] = \Pr[\mathcal{B}_j = 1] = \epsilon \ge \epsilon - \epsilon^2$.

For the induction step, assume that both the upper and lower bounds hold upto $t-1$ variables. Let $p_t := \Pr[\mathcal{B}_t = 1]$ and also denote the $(t-1)$-sum by $\mathcal{S}' = \sum_{j=1}^{t-1} \mathcal{B}_j$. Then we have

$$\begin{aligned}
\Pr\left[\mathcal{S} \texttt{ is odd}\right] &= \Pr\left[\mathcal{S}' \texttt{ is even and } \mathcal{B}_t = 1\right] + \Pr\left[\mathcal{S}' \texttt{ is odd and } \mathcal{B}_t = 0\right] \\
&\ge (1 - (\epsilon - p_t))p_t + ((\epsilon - p_t) - (\epsilon - p_t)^2)(1 - p_t) \quad \text{(By the inductive hypothesis)} \\
&= p_t - \epsilon p_t + p_t^2 + \left(\epsilon - p_t - \epsilon^2 + 2\epsilon p_t - p_t^2\right) - \left((\epsilon - p_t) - (\epsilon - p_t)^2\right)p_t \\
&= \epsilon - \epsilon^2 + \left(1 - \epsilon - 1 + 2\epsilon - \epsilon + (\epsilon - p_t)^2\right)p_t + (1 - 1 + 1)p_t^2 \\
&= \epsilon - \epsilon^2 + (\epsilon - p_t)^2 p_t + p_t^2 \ \ge \epsilon - \epsilon^2.
\end{aligned}$$

Hence, the lower bound holds also for t. □