Guessing Attacks on Distributed-Storage Systems

Annina Bracher
ETH Zurich

Eran Hof
R&D Center, Ramat-Gan

Amos Lapidoth
ETH Zurich

***Abstract*—We study the secrecy of a distributed-storage system for passwords. The encoder, Alice, observes a length-n password and describes it using δ s-bit hints, which she stores in different locations. The legitimate receiver, Bob, observes ν of those hints. In one scenario we require that the expected number of guesses it takes Bob to guess the password approach 1 as n tends to infinity, and in the other that the expected size of the shortest list that Bob must form to guarantee that it contain the password approach 1. The eavesdropper, Eve, sees $\eta < \nu$ hints. Assuming that Alice cannot control which hints Bob and Eve observe, we characterize for each scenario the largest normalized (by n) exponent that we can guarantee for the expected number of guesses it takes Eve to guess the password.**

I. Introduction

We generalize the model that was studied in [1] to allow for Alice to produce δ hints (not necessarily 2), for Bob to observe η (not necessarily 2) of those hints, and for Eve to observe $\nu < \eta$ (not necessarily 1) of the hints. We thus consider the following scenario: Some sensitive information X (e.g. password) is drawn from its finite support $\mathcal{X}$ according to some PMF P_X. A (stochastic) encoder, Alice, maps (possibly using randomization) X to a set of hints and stores them on different discs in different locations. The hints are intended for a legitimate receiver, Bob, who knows where they are stored and thus sees more hints than an eavesdropper, Eve. Given some notion of ambiguity, we would ideally like Bob's ambiguity about X to be small and Eve's large.

We require that the network be robust against a limited number of disc failures: unlike the model in [1] where Alice produces two hints, Bob sees both hints, and Eve sees one hint, here we assume that Alice produces δ s-bit hints $M_1, M_2, \ldots, M_\delta$, Bob sees $\nu \leq \delta$ hints, and Eve sees $\eta < \nu$ hints. Which hints Bob and Eve observe is a subtle question. We adopt a conservative approach and assume that, after observing X, an adversarial genie reveals to Bob the ν hints that maximize his ambiguity and to Eve the η hints that minimize her ambiguity. Not allowing the genie to observe X would lead to a weaker form of secrecy (see Example 1).

The considered network is a distributed-storage system, which is static in the sense that failed discs are not replaced. The case where X is drawn uniformly, Bob must reconstruct X, and Eve's observation must satisfy some information-theoretic security criterion (e.g., the mutual information between Eve's observation and X must be null) corresponds to the erasure-erasure wiretap channel studied in [2] and is a special case of the wiretap networks in [3], [4]. In the literature, the setting is also known as "secret sharing". In traditional secret sharing, each set of hints either reveals X or reveals no information about X [5], [6]. More general are ramp schemes, where any ν hints reveal X and the amount of information that fewer-than-ν hints reveal is controlled (see e.g. [7]). Our setting is different in that we assume $X \sim P_X$ and in that, using some notion of ambiguity, we quantify how difficult it is for Bob and Eve to reconstruct X.

There are several ways to define ambiguity. For example, we could require that Bob be able to reconstruct X whenever X is "typical" and that the conditional entropy of X given Eve's observation be large. For some scenarios, such an approach might be inadequate. Firstly, this approach may not properly address Bob's needs when X is not typical. For example, if Bob must guess X, this approach does not guarantee that the expected number of guesses be small: it only guarantees that the probability of success after one guess be large. It does not indicate the number of guesses that Bob might need when X is atypical. Secondly, conditional entropy need not be an adequate measure of Eve's ambiguity: if Eve tries to guess X, then we may care more about the number of guesses that Eve needs than about the conditional entropy [8].

In this paper, we assume that Eve wants to guess X with the least number of guesses of the form "Is X = x?". We quantify Eve's ambiguity about X by the expected number of guesses that she needs to uncover X. In this sense, Eve faces an instance of the Massey-Arikan guessing problem [9]. For each possible observation z in some finite set $\mathcal{Z}$, Eve chooses a guessing function $G(\cdot\,|z)$ from $\mathcal{X}$ onto the set $\{1, \ldots, |\mathcal{X}|\}$, which determines the guessing order: if Eve observes z, then "Is $X = x$?" will be her $G(x|z)$-th question. Eve's expected number of guesses is $\mathbb{E}[G(X|Z)]$. This expectation is minimized if for each $z \in \mathcal{Z}$ the guessing function $G(\cdot\,|z)$ orders the possible realizations of X in decreasing order of their posterior probabilities given $Z = z$.

As to Bob, we consider two different criteria: In the "guessing version" the criterion is the expected number of guesses it takes Bob to guess X, and in the "list version" the criterion is the first moment of the size of the list that Bob must form to guarantee that it contain X. The merits of the two criteria are discussed in [1]. They lead to similar results (see Section IV).

The list-size criterion can be viewed as a worst-case version of the guessing criterion: If Bob tries to guess X without knowing its PMF, then it is reasonable for him to first guess the possible realizations of X that Alice could have mapped to the observed hints. The number of guesses that he needs is then at most the size of the smallest list that is guaranteed to contain X. Section V elaborates on why we allow Eve to

guess also in the list version.

With no extra effort we can generalize the model and replace expectations with ρ-th moments. This we do to better bring out the role of Rényi entropy. For an arbitrary $\rho > 0$, we thus study the ρ-th (instead of the first) moment of the list-size and of the number of guesses. Moreover, we allow some side-information Y that is available to all parties. We thus assume that the pair (X, Y) takes value in the finite set $\mathcal{X} \times \mathcal{Y}$ according to $P_{X,Y}$.

II. Problem Statement

We consider two problems: the "guessing version" and the "list version". They differ in the definition of Bob's ambiguity. In both versions a pair (X, Y) is drawn from the finite set $\mathcal{X} \times \mathcal{Y}$ according to the PMF $P_{X,Y}$, and $\rho > 0$ is fixed. Denote by $\mathbb{F}_q$ the Galois field with q elements. Upon observing $(X, Y) = (x, y)$, Alice draws the δ-tuple $\boldsymbol{M} = (M_1, \ldots, M_\delta)$ from the finite set $\mathbb{F}_{2^s}^\delta$ according to some conditional PMF

$$\mathbb{P}[\boldsymbol{M} = \mathbf{m} \,|X = x, Y = y\,], \, \mathbf{m} \in \mathbb{F}_{2^s}^\delta. \tag{1}$$

An adversary chooses a size-ν set $\mathcal{B} \subseteq \{1, \ldots, \delta\}$ and reveals to Bob the set $\mathcal{B}$ and the components $\boldsymbol{M}_\mathcal{B}$ of $\boldsymbol{M}$ indexed by $\mathcal{B}$. Based on $\boldsymbol{M}_\mathcal{B}$ and the side-information Y, Bob guesses X using an optimal guessing function $G_\mathcal{B}$, which minimizes the ρ-th moment of the number of guesses that he needs. As indicated by the subscript, $G_\mathcal{B}$ depends on $\mathcal{B}$. In the "guessing version" we define Bob's min-max ambiguity about X as

$$\mathscr{A}_{\mathrm{B}}^{(\mathrm{g})}(P_{X,Y}) = \min_{G_\mathcal{B}} \mathbb{E}\Big[\max_{\mathcal{B}} G_\mathcal{B}(X|Y, \boldsymbol{M}_\mathcal{B})^\rho\Big], \tag{2}$$

and in the "list version" as

$$\mathscr{A}_{\mathrm{B}}^{(\mathrm{l})}(P_{X,Y}) = \mathbb{E}\Big[\max_{\mathcal{B}} |\mathcal{L}^Y_{\boldsymbol{M}_\mathcal{B}}|^\rho\Big], \tag{3}$$

where for all $y \in \mathcal{Y}$ and $\mathbf{m}_\mathcal{B} \in \mathbb{F}_{2^s}^\delta$

$$\mathcal{L}^y_{\mathbf{m}_\mathcal{B}} = \big\{x \colon \mathbb{P}[X = x \,|Y = y, \boldsymbol{M}_\mathcal{B} = \mathbf{m}_\mathcal{B}\,] > 0\big\}. \tag{4}$$

Eve gets to see a size-η set $\mathcal{E} \subseteq \{1, \ldots, \delta\}$ and the components $\boldsymbol{M}_\mathcal{E}$ of $\boldsymbol{M}$ indexed by $\mathcal{E}$. The set is chosen by an accomplice of hers. Based on $\boldsymbol{M}_\mathcal{E}$ and the side-information Y, Eve guesses X. In both versions we define her ambiguity about X as

$$\mathscr{A}_{\mathrm{E}}(P_{X,Y}) = \min_{G_\mathcal{E}} \mathbb{E}\Big[\min_{\mathcal{E}} G_\mathcal{E}(X\,|Y, \boldsymbol{M}_\mathcal{E}\,)^\rho\Big]. \tag{5}$$

Optimizing over Alice's choice of the conditional PMF in (1), we wish to characterize the largest ambiguity that we can guarantee that Eve will have subject to a given upper bound on the ambiguity that Bob may have.

Of special interest to us is the asymptotic regime where (X, Y) is an n-tuple (not necessarily drawn IID), and where each hint stores

$$s = nR_s$$

bits, where R_s is nonnegative and corresponds to the per-hint storage-rate. (We assume that δ, ν, and η are fixed.) We characterize the largest exponential growth for Eve's ambiguity that can be guaranteed subject to Bob's ambiguity tending to 1. This asymptote turns out not to depend on the version, and in the asymptotic analysis $\mathscr{A}_{\mathrm{B}}$ can stand for either $\mathscr{A}_{\mathrm{B}}^{(\mathrm{g})}$ or $\mathscr{A}_{\mathrm{B}}^{(\mathrm{l})}$.

To phrase this mathematically, let us introduce the stochastic process $\{(X_i, Y_i)\}_{i\in\mathbb{N}}$ with finite alphabet $\mathcal{X} \times \mathcal{Y}$. We denote by P_{X^n,Y^n} the PMF of (X^n, Y^n). For a nonnegative rate R_s, we call E_{E} an *achievable ambiguity-exponent* if there is a sequence of stochastic encoders such that Bob's ambiguity (which is always at least 1) satisfies

$$\lim_{n\to\infty} \mathscr{A}_{\mathrm{B}}(P_{X^n,Y^n}) = 1, \tag{6}$$

and such that Eve's ambiguity satisfies

$$\liminf_{n\to\infty} \frac{\log(\mathscr{A}_{\mathrm{E}}(P_{X^n,Y^n}))}{n} \geq E_{\mathrm{E}}. \tag{7}$$

We characterize the supremum $\overline{E_{\mathrm{E}}}$ of all achievable ambiguity-exponents, which we call *privacy-exponent*. If (6) cannot be satisfied, then the set of achievable ambiguity-exponents is empty, and we say that the privacy-exponent is negative infinity.

III. Main Results

To describe our results, we shall need the conditional Rényi entropy [9]–[11]

$$H_\alpha(X|Y) = \frac{\alpha}{1-\alpha} \log \sum_{y\in\mathcal{Y}} \Big(\sum_{x\in\mathcal{X}} P_{X,Y}(x,y)^\alpha\Big)^{1/\alpha}, \tag{8}$$

where $\alpha \in [0, \infty]$ is the order and where the cases where α is 0, 1, or ∞ are treated by limiting arguments. In addition, we shall need the notion of conditional Rényi entropy-rate: Let $\{(X_i, Y_i)\}_{i\in\mathbb{N}}$ be a discrete-time stochastic process with finite alphabet $\mathcal{X} \times \mathcal{Y}$. Whenever the limit as n tends to infinity of $H_\alpha(X^n|Y^n)\,/n$ exists, we denote it by $H_\alpha(\boldsymbol{X}|\boldsymbol{Y})$ and call it conditional Rényi entropy-rate. In this paper, $\alpha = 1/(1+\rho)$ takes value in the set $(0, 1)$. To simplify notation, we henceforth write $\tilde{\rho}$ for $1/(1+\rho)$. We denote by $\alpha \vee \beta$ the maximum of α and β and by $\alpha \wedge \beta$ their minimum.

A. Finite Blocklength Results

In the next two theorems $(\nu - \eta)r$ should be viewed as the number of information-bits that can be gleaned about X from ν but not from η hints and γp as the number of information-bits that any $\gamma \leq \nu$ hints reveal about X (see Section VI).

Theorem 1 (Finite Blocklength Guessing Version): For $p, r \in \{0, 1, \ldots, s\}$ satisfying

$$p + r = s \tag{9a}$$

$$p, r \in \{0\} \cup \big\{\lceil \log \delta \rceil, \lceil \log \delta \rceil + 1, \ldots\big\}, \tag{9b}$$

there is a choice of the conditional PMF in (1) for which Bob's ambiguity about X is upper-bounded by

$$\mathscr{A}_{\mathrm{B}}^{(\mathrm{g})}(P_{X,Y}) < 1 + 2^{\rho(H_{\tilde{\rho}}(X|Y) - \nu s + \eta r + 1)}, \tag{10}$$

and Eve's ambiguity about X is lower-bounded by

$$\mathscr{A}_{\mathrm{E}}(P_{X,Y}) \geq 2^{\rho(H_{\tilde{\rho}}(X|Y) - \eta(s-r) - \eta \log \delta - \log(1+\ln|\mathcal{X}|))}. \tag{11}$$

Conversely, for every conditional PMF, Bob's ambiguity is lower-bounded by

$$\mathscr{A}_{\mathrm{B}}^{(\mathrm{g})}(P_{X,Y}) \geq 2^{\rho(H_{\tilde{\rho}}(X|Y) - \nu s - \log(1+\ln|\mathcal{X}|))} \vee 1, \tag{12}$$

and Eve's ambiguity is upper-bounded by

$$\mathscr{A}_{\mathrm{E}}(P_{X,Y}) \leq 2^{\rho(\nu-\eta)s}\mathscr{A}_{\mathrm{B}}^{(\mathrm{g})}(P_{X,Y}) \wedge 2^{\rho H_{\tilde{p}}(X|Y)}. \tag{13}$$

Theorem 2 (Finite Blocklength List Version): Suppose that $2^{\nu s} > \log|\mathcal{X}| + 2$. For p, $r \in \{0, 1, \ldots, s\}$ satisfying (9) and

$$2^{\nu s - \eta r} > \log|\mathcal{X}| + 2, \tag{14a}$$

there is a choice of the conditional PMF in (1) for which Bob's ambiguity about X is upper-bounded by

$$\mathscr{A}_{\mathrm{B}}^{(\mathrm{l})}(P_{X,Y}) < 1 + 2^{\rho(H_{\tilde{p}}(X|Y) - \log(2^{\nu s - \eta r} - \log|\mathcal{X}| - 2) + 2)}, \tag{15}$$

and Eve's ambiguity about X is lower-bounded by

$$\mathscr{A}_{\mathrm{E}}(P_{X,Y}) \geq 2^{\rho(H_{\tilde{p}}(X|Y) - \eta(s-r) - \eta\log\delta - \log(1+\ln|\mathcal{X}|))}. \tag{16}$$

Conversely, for every conditional PMF, Bob's ambiguity is lower-bounded by

$$\mathscr{A}_{\mathrm{B}}^{(\mathrm{l})}(P_{X,Y}) \geq 2^{\rho(H_{\tilde{p}}(X|Y) - \nu s)} \vee 1, \tag{17}$$

and Eve's ambiguity is upper-bounded by

$$\mathscr{A}_{\mathrm{E}}(P_{X,Y}) \leq 2^{\rho(\nu-\eta)s}\mathscr{A}_{\mathrm{B}}^{(\mathrm{l})}(P_{X,Y}) \wedge 2^{\rho H_{\tilde{p}}(X|Y)}. \tag{18}$$

Proof: The proof is outlined in Section VI. ∎

The bounds in Theorems 1 and 2 are tight in the sense that with a judicious choice of p and r the achievability results (namely (10)–(11) in the "guessing version" and (15)–(16) in the "list version") match the corresponding converse results (namely (12)–(13) in the "guessing version" and (17)–(18) in the "list version") up to polynomial factors of δ^{η} and of $\ln|\mathcal{X}|$. This can be seen from the following corollary:

Corollary 3: In the guessing version, for any

$$\mathscr{U}_{\mathrm{B}} \geq 1 + 2^{\rho(H_{\tilde{p}}(X|Y) - \nu s + 1)}, \tag{19}$$

there is a choice of the conditional PMF in (1) for which Bob's ambiguity about X is upper-bounded by

$$\mathscr{A}_{\mathrm{B}}^{(\mathrm{g})}(P_{X,Y}) < \mathscr{U}_{\mathrm{B}}, \tag{20}$$

and Eve's ambiguity about X is lower-bounded by

$$\begin{aligned}\mathscr{A}_{\mathrm{E}}(P_{X,Y}) \,/\big[\delta^{\eta}(1+\ln|\mathcal{X}|)\big]^{-\rho} \\ \geq \Big(\big[(2\delta)^{-\rho\eta}2^{\rho(\nu-\eta)s}(\mathscr{U}_{\mathrm{B}} - 1)\big] \wedge 2^{\rho H_{\tilde{p}}(X|Y)}\Big).\end{aligned} \tag{21}$$

In the list version, for any

$$\mathscr{U}_{\mathrm{B}} \geq 1 + 2^{\rho(H_{\tilde{p}}(X|Y) - \log(2^{\nu s} - \log|\mathcal{X}| - 2) + 2)}, \tag{22}$$

there is a choice of the conditional PMF in (1) for which Bob's ambiguity about X is upper-bounded by

$$\mathscr{A}_{\mathrm{B}}^{(\mathrm{l})}(P_{X,Y}) < \mathscr{U}_{\mathrm{B}}, \tag{23}$$

and Eve's ambiguity about X is lower-bounded by

$$\begin{aligned}\mathscr{A}_{\mathrm{E}}(P_{X,Y}) \,/\big[\delta^{\eta}(1+\ln|\mathcal{X}|)\big]^{-\rho} \\ \geq \Big(\big[2^{-3\rho}(2\delta)^{-\rho\eta}2^{\rho(\nu-\eta)s}(\mathscr{U}_{\mathrm{B}} - 1)\big] \wedge 2^{\rho H_{\tilde{p}}(X|Y)} \\ \wedge \big[\{2(2\delta)^{\eta}(2+\log|\mathcal{X}|)\}^{-\rho}2^{\rho((\nu-\eta)s + H_{\tilde{p}}(X|Y))}\big]\Big).\end{aligned} \tag{24}$$

We conclude the section with some remarks. First, we want to understand why it is a good idea to store an equal number of bits on each disc. This can be seen from the next result:

Theorem 4: Suppose that for $\ell \in \{1, \ldots, \delta\}$ Disc ℓ stores s_ℓ bits, where $s_1 \leq \ldots \leq s_\delta$. Then, depending on the version of the problem, Bob's ambiguity is for every conditional PMF lower-bounded by

$$\mathscr{A}_{\mathrm{B}}^{(\mathrm{g})}(P_{X,Y}) \geq 2^{\rho\left(H_{\tilde{p}}(X|Y) - \sum_{\ell=1}^{\nu} s_\ell - \log(1+\ln|\mathcal{X}|)\right)} \vee 1 \tag{25}$$

$$\mathscr{A}_{\mathrm{B}}^{(\mathrm{l})}(P_{X,Y}) \geq 2^{\rho\left(H_{\tilde{p}}(X|Y) - \sum_{\ell=1}^{\nu} s_\ell\right)} \vee 1, \tag{26}$$

and Eve's ambiguity is upper-bounded by

$$\mathscr{A}_{\mathrm{E}}(P_{X,Y}) \leq 2^{\rho\sum_{\ell=1}^{\nu-\eta} s_\ell}\mathscr{A}_{\mathrm{B}}^{(\mathrm{g})}(P_{X,Y}) \wedge 2^{\rho H_{\tilde{p}}(X|Y)} \tag{27}$$

$$\mathscr{A}_{\mathrm{E}}(P_{X,Y}) \leq 2^{\rho\sum_{\ell=1}^{\nu-\eta} s_\ell}\mathscr{A}_{\mathrm{B}}^{(\mathrm{l})}(P_{X,Y}) \wedge 2^{\rho H_{\tilde{p}}(X|Y)}. \tag{28}$$

The theorem and Corollary 3 imply the following:

*Note 5 (**Why store s bits on each disc?**): Compare a scenario where for $\ell \in \{1, \ldots, \delta\}$ Disc ℓ stores s_ℓ bits with one where each disc stores $\lfloor (s_1 + \ldots + s_\delta)/\delta \rfloor$ bits. Neglecting polynomial factors of δ^{η} and of $\ln|\mathcal{X}|$, each pair of ambiguities for Bob and Eve that is achievable in the former scenario is also achievable in the latter scenario.*

Next, we explain why we did not define Eve's ambiguity as

$$\tilde{\mathscr{A}_{\mathrm{E}}}(P_{X,Y}) = \min_{\mathcal{E}} \min_{G_{\mathcal{E}}} \mathbb{E}\big[G_{\mathcal{E}}(X\,|Y, \boldsymbol{M}_{\mathcal{E}})^{\rho}\big]. \tag{29}$$

*Note 6 (**Which hints does Eve observe?**): Quantifying Eve's ambiguity by* (5)*, we assume that after observing $(X, Y, \boldsymbol{M})$ an adversarial "genie" reveals to Eve the hints that minimize her guessing efforts. Less conservative is* (29)*, which applies if Eve observes the hints that in expectation minimize her guessing efforts. If Eve's ambiguity were quantified by* (29)*, then Theorems 1 and 2 would still apply.*[1] *However,* (29) *leads to a weaker form of secrecy than* (5)*.*

We illustrate this by means of an example:

Example 1: Suppose that Y is null and that X is drawn uniformly from $\mathcal{X}$. Take $\delta = \nu = 2$ and $\eta = 1$. Let Alice describe X using $(M_1 = X, M_2 = *)$ or $(M_1 = *, M_2 = X)$, where $*$ is not in $\mathcal{X}$, each with probability $1/2$. Since Bob can recover X from (M_1, M_2), we have $\mathscr{A}_B^{(g)}(P_{X,Y}) = \mathscr{A}_B^{(l)}(P_{X,Y}) = 1$. The probability that $M_1 = *$ is $1/2$ and likewise for M_2. Thus, regardless of whether $\mathcal{E} = 1$ or $\mathcal{E} = 2$, if Eve always observes the same hint, then the ρ-th moment of the number of guesses that she needs is at least $\min_G \mathbb{E}[G(X)^{\rho}]/2$ and $\tilde{\mathscr{A}_E}(P_{X,Y}) \geq \min_G \mathbb{E}[G(X)^{\rho}]/2$. However, one of the two hints always reveals X, and therefore we have $\mathscr{A}_E(P_{X,Y}) = 1$.

B. Asymptotic Results

Consider now the asymptotic regime where (X, Y) is an n-tuple, and where n is large. In this case the results are the same for both versions, and we thus refer to both $\mathscr{A}_{\mathrm{B}}^{(\mathrm{g})}$ and $\mathscr{A}_{\mathrm{B}}^{(\mathrm{l})}$ by $\mathscr{A}_{\mathrm{B}}$. Theorems 1 and 2 and Corollary 3 imply the following asymptotic result:

[1] We could even tighten (11) and (16): under (29) the subtraction of $\rho\eta\log\delta$ in the exponents of the lower bounds (11) and (16) is not needed since the genie cannot use its choice of $\mathcal{E}$ to convey to Eve information about X.

Corollary 7: Let $\{(X_i, Y_i)\}_{i\in\mathbb{N}}$ be a discrete-time stochastic process with finite alphabet $\mathcal{X}\times\mathcal{Y}$, and suppose its conditional Rényi entropy-rate $H_{\tilde{\rho}}(\boldsymbol{X}|\boldsymbol{Y})$ is well-defined. Given any positive rate R_s, the privacy-exponent is

$$\overline{E_{\mathrm{E}}} = \begin{cases} \rho\big(R_s(\nu-\eta) \wedge H_{\tilde{\rho}}(\boldsymbol{X}|\boldsymbol{Y})\big), & \nu R_s > H_{\tilde{\rho}}(\boldsymbol{X}|\boldsymbol{Y}) \\ -\infty, & \nu R_s < H_{\tilde{\rho}}(\boldsymbol{X}|\boldsymbol{Y})\,. \end{cases} \tag{30}$$

To achieve the maximal privacy-exponent $\rho H_{\tilde{\rho}}(\boldsymbol{X}|\boldsymbol{Y})$, the per-hint storage-rate must satisfy $R_s \geq H_{\tilde{\rho}}(\boldsymbol{X}|\boldsymbol{Y})/(\nu-\eta)$, where $H_{\tilde{\rho}}(\boldsymbol{X}|\boldsymbol{Y})$ is the minimum rate that is necessary to describe the source for Bob. This agrees with the well-known result that the optimal share-size to share a k-bit secret so that any ν shares reveal X and any η shares provide no information about X is $k/(\nu-\eta)$ (see e.g. [2]).

IV. Lists and Guesses

Suppose that a guessing decoder and a list-decoder try to retrieve X from the side-information Y and some description Z, which an encoder produces after observing (X,Y). To prove Theorems 1 and 2, we must understand how small the decoders' ambiguities about X can be and how they relate to each other. This section resolves the above issues.

In the following, $G^*(\cdot|Y)$ denotes a guessing function that minimizes $\mathbb{E}[G(X|Y)^\rho]$, and the lists $\{\mathcal{L}^y\}$ stand for the sets

$$\mathcal{L}^y = \{x\colon P_{X|Y}(x|y) > 0\},\ y\in\mathcal{Y}. \tag{31}$$

Denote the finite support of Z by $\mathcal{Z}$. For some conditional PMF $\mathbb{P}[Z=z|X=x,Y=y]$, $z\in\mathcal{Z}$, $(x,y)\in\mathcal{X}\times\mathcal{Y}$, denote by $G^*(\cdot|Y,Z)$ a guessing function that minimizes $\mathbb{E}\big[G(X|Y,Z)^\rho\big]$, and define the lists

$$\mathcal{L}_z^y = \{x\colon \mathbb{P}[X=x|Y=y,Z=z]>0\},\ y\in\mathcal{Y},\ z\in\mathcal{Z}.$$

If $|\mathcal{Z}| > \log|\mathcal{X}|+2$, then an upper and a lower bound on the smallest ambiguity $\mathbb{E}\big[|\mathcal{L}_Z^Y|^\rho\big]$ of a list-decoder, which are tight up to polylogarithmic factors of $|\mathcal{X}|$, are given in [10, Theorem VI.1]. If Z is null, then an upper and a lower bound on the ambiguity $\min\mathbb{E}[G(X|Y)^\rho]$ of a guessing decoder can be found in [9, Theorem 1 and Proposition 4]. The following result quantifies how additional side-information Z helps guessing: it shows that Z can reduce the ρ-th moment of the number of guesses by at most a factor of $|\mathcal{Z}|^{-\rho}$.

Corollary 8: [1, Corollary 10] Let (X,Y) be drawn from the finite set $\mathcal{X}\times\mathcal{Y}$ according to the PMF $P_{X,Y}$, and let $\mathcal{Z}$ be a finite set. There exists a function $f\colon\mathcal{X}\times\mathcal{Y}\to\mathcal{Z}$ such that for $Z=f(X,Y)$

$$\min_G \mathbb{E}[G(X|Y,Z)^\rho] < 1 + 2^\rho\,|\mathcal{Z}|^{-\rho}\min_G \mathbb{E}[G(X|Y)^\rho]\,. \tag{32}$$

Conversely, for every chance variable Z with alphabet $\mathcal{Z}$

$$\min_G \mathbb{E}[G(X|Y,Z)^\rho] \geq |\mathcal{Z}|^{-\rho}\min_G \mathbb{E}\big[G(X|Y)^\rho\big] \vee 1. \tag{33}$$

From Corollary 8 and [9, Theorem 1 and Proposition 4] we obtain an upper and a lower bound on the smallest ambiguity $\min_G\mathbb{E}[G(X|Y,Z)^\rho]$ that is achievable for a given $|\mathcal{Z}|$. The bounds are tight up to polylogarithmic factors of $|\mathcal{X}|$.

We next provide a link between guessing and list-decoding:

Theorem 9: [1, Special Case of Theorem 5] Let (X,Y) be drawn from the finite set $\mathcal{X}\times\mathcal{Y}$ according to the PMF $P_{X,Y}$, and let $Z=\lfloor\log G^*(X|Y)\rfloor$. Then, $\mathcal{Z}=\{0,\ldots,\lfloor\log|\mathcal{X}|\rfloor\}$ and the following hold:

$$\mathbb{P}\big[G^*(X|Y)\leq|\mathcal{L}^Y|\big]=1, \quad |\mathcal{L}_Z^Y|\leq G^*(X|Y)\,. \tag{34}$$

On account of Corollary 8, providing $Z=\lfloor\log G^*(X|Y)\rfloor$ to a guessing decoder can improve its performance by at most a polynomial factor of $|\mathcal{Z}|=\lfloor\log|\mathcal{X}|\rfloor+1$. Hence, Theorem 9 explains why the results for the "guessing version" and the "list version" differ only by polylogarithmic factors of $|\mathcal{X}|$.

Theorem 9 is especially interesting in asymptotic settings, where polylogarithmic factors of $|\mathcal{X}|$ are negligible. For example, one can combine Theorem 9 and Corollary 8 with [9, Theorem 1 and Proposition 4] to provide upper and lower bounds on the smallest $\mathbb{E}\big[|\mathcal{L}_Z^Y|^\rho\big]$ that is achievable for a given $|\mathcal{Z}|$. These bounds are weaker (by at most polylogarithmic factors of $|\mathcal{X}|$) than [10, Theorems I.1 and VI.1] in the finite blocklength regime but tight enough to prove the asymptotic results [10, Theorems I.2 and VI.2].

Another example where Theorem 9 is useful is in determining the feedback listsize capacity of a DMC $W(y|x)$ with positive zero-error capacity: Theorem 9 can be used to give an elegant proof of the direct part of [13, Theorem I.1], which states that in the presence of perfect feedback the listsize capacity of $W(y|x)$ equals the cutoff rate $R_{\text{cutoff}}(\rho)$ with feedback (which is in fact equal to the cutoff rate without feedback [13, Corollary I.4]). To see this, suppose that we are given (feedback) codes of rate R for which the ρ-th moment of the number of guesses $G^*(M|Y^n)$ a decoder needs to guess the transmitted message M based on the channel-outputs Y^n approaches 1 as the blocklength n tends to infinity. (Recall that $R_{\text{cutoff}}(\rho)$ is the supremum of all rates for which such codes exist.) Suppose now that the transmission does not stop after n channel uses. Instead, the encoder computes $Z=\lfloor\log G^*(M|Y^n)\rfloor\in\{0,\ldots,\lfloor nR\rfloor\}$ from the feedback Y^n and uses another n' channel uses to transmit Z at a positive rate while guaranteeing that the receiver can decode it with probability 1. Since a positive zero-error (feedback) capacity cannot be smaller than 1, it is enough to take $n'\leq\lceil\log(nR)\rceil$. Hence, $(n+n')/n$ converges to 1 as n tends to infinity, and the rate of the code thus converges to R. At the same time, Theorem 9 implies that the size of the smallest decoding-list $\mathcal{L}^{Y^{n+n'}}$ that is guaranteed to contain M satisfies $|\mathcal{L}^{Y^{n+n'}}|=|\mathcal{L}_Z^{Y^n}|\leq G^*(M|Y^n)$, and thus the ρ-th moment of $|\mathcal{L}^{Y^{n+n'}}|$ converges to 1 as n tends to infinity. This proves that in the presence of perfect feedback the listsize capacity of $W(y|x)$ is lower-bounded by $R_{\text{cutoff}}(\rho)$.

V. A Fair Opponent for a List-Decoder

This section elaborates on why in the list version Eve does not form lists but guesses X. We first argue that a list-decoder is not a fair opponent for a list-decoder:

Example 2: Suppose that Y is null and that X is an n-tuple X^n, where $\{X_i\}_{i\in\mathbb{N}}$ are IID $\mathrm{Ber}(1/2)$. *Define* $\{Z_i\}_{i\in\mathbb{N}}$

by $Z_i = X_i \oplus S_i$, where $\{S_i\}_{i\in\mathbb{N}}$ are IID $\mathrm{Ber}(\epsilon)$ for some $\epsilon \in (0, 1/2)$. Note that because $\mathbb{P}[X^n = \mathbf{x}|Z^n = \mathbf{z}] > 0$, $\forall\, \mathbf{x}, \mathbf{z} \in \{0,1\}^n$ we have $\mathbb{E}[|\mathcal{L}_{Z^n}|^\rho] = 2^n$, where for all $\mathbf{z} \in \{0,1\}^n$

$$\mathcal{L}_{\mathbf{z}} = \{\mathbf{x} \colon \mathbb{P}[X^n = \mathbf{x}|Z^n = \mathbf{z}, W_n = w_n] > 0\}.$$

[10, Theorem VI.2] implies that for $R > H_{\tilde{\rho}}(X|Z)$ there exist $W_n \in \{1, \dots, 2^{nR}\}$ s.t. $\lim_{n\to\infty} \mathbb{E}[|\mathcal{L}_{Z^n,W_n}|^\rho] = 1$. Suppose now that Bob can recover (Z^n, W_n) and Eve Z^n, but that conditional on Z^n Eve's observation is independent of W_n. If both Bob and Eve must form lists that are guaranteed to contain X^n, then the ρ-th moment of the size of Bob's list tends to 1 as n increases, and the ρ-th moment of the size of Eve's list is $2^{\rho n}$ and hence as large as it can be. This can be used to show that if Eve must form a list that is guaranteed to contain X, then secrecy comes almost for free: Indeed, the minimum required rate to describe X^n so that the ρ-th moment of the size of Bob's list tends to 1 is 1 (see [10, Theorem VI.2]). Because $\lim_{\epsilon\downarrow 0} H_{\tilde{\rho}}(X|Z) = 0$ and because rate 1 suffices to describe Z^n perfectly, we can achieve rates arbitrarily close to 1 by describing X by (Z^n, W_n). Moreover, perfect secrecy is attained if Eve sees only Z^n but not W_n because in this case the ρ-th moment of the size of Eve's list is $2^{\rho n}$ and thus as large as it can be. But encrypting W_n is cheap because W_n takes value in a set of size 2^{nR}, where R can be arbitrarily close to 0 for ϵ sufficiently small. In contrast, the ρ-th moment of the number of guesses that Eve needs to guess X is at most $2^{n\rho R}$ because $H_{\tilde{\rho}}(X^n|Z^n) < nR$ (see [9, Proposition 4]).

In the classical Shannon cipher system [14], a popular way to measure imperfect security is in terms of equivocation, i.e., in terms of the conditional entropy $H(X|Z)$ of the sensitive information X given Eve's observation Z. In the setting where Bob is a list-decoder or a guessing decoder, Rényi entropy plays the role of Shannon entropy in the sense that the minimum required rate to encode an n-tuple $X = X^n$ is the Rényi entropy rate $H_{\tilde{\rho}}(\boldsymbol{X})$ rather than the Shannon entropy rate $H(\boldsymbol{X}) = H_1(\boldsymbol{X})$. Therefore, the conditional Rényi entropy $H_{\tilde{\rho}}(X|Z)$ qualifies as a natural equivalent for equivocation. Note that $H_{\tilde{\rho}}(X|Z)$ has a nice operational characterization: $2^{\rho H_{\tilde{\rho}}(X|Z)}$ is (up to polylogarithmic factors of the size of the support of X) the ρ-th moment of the number of guesses that Eve needs to guess X [9].

VI. How to Prove Theorems 1 and 2

The converse results ((12)–(13) in the "guessing version" and (17)–(18) in the "list version") readily follow from the results in Section IV. To prove the achievability results ((10)–(11) in the "guessing version" and (15)–(16) in the "list version"), we use the following coding scheme. Upon observing (X, Y), Alice describes X deterministically by a tuple (V, W), where V takes value in the finite field $\mathbb{F}_{2^p}^{\nu}$ and W in $\mathbb{F}_{2^r}^{\nu-\eta}$. Depending on the version, she chooses the description so that if Bob's observation were (V, W), then his ambiguity about X would satisfy the upper bound (10) in the "guessing version" or (15) in the "list version". Then, she maps V to a length-δ codeword of a $(\delta, \nu, \delta - \nu + 1)$ maximum distance separable (MDS) code over $\mathbb{F}_{2^p}$ and stores each codeword symbol on a different disc. Since the code is MDS, any $\gamma \le \nu$ hints reveal γp bits of V. Independently of (X, Y), Alice draws a random variable U uniformly from the field $\mathbb{F}_{2^r}^{\eta}$, maps (W, U) to a length-δ codeword of a $(\delta, \nu, \delta - \nu + 1)$ MDS code over the field $\mathbb{F}_{2^r}$, and stores each codeword symbol on a different disc. She chooses the mapping so that any η codeword symbols are independent of W or, equivalently, that given W it is possible to reconstruct U from any η codeword symbols. (As in [2], this is accomplished using nested MDS codes.) As a consequence, W can be recovered from any ν hints while any η hints reveal no information about W.

Summing up, the outlined coding scheme guarantees that, after observing ν hints, Bob can reconstruct the tuple (V, W). Hence, his ambiguity about X satisfies (10) in the "guessing version" and (15) in the "list version". Observing η hints enables Eve to recover ηp bits of V, but it does not enable her to recover any information about W. Using the results of Section IV, we can thus show that observing η hints can decrease Eve's guessing efforts by at most a factor of $2^{-\rho\nu p}$.[2] Since we quantify Eve's ambiguity by (5), we assume that upon observing all the hints and (X, Y) an adversarial genie reveals to Eve the η hints that minimize her ambiguity (see Note 6). In doing so, the genie can decrease Eve's ambiguity by an additional factor of at most $\delta^{-\rho\eta}$ (this is due to Corollary 8 and the fact that there are $\binom{\delta}{\eta} \le \delta^\eta$ size-η subsets of $\{1, \dots, \delta\}$).

References

[1] A. Bracher, E. Hof, and A. Lapidoth, "Distributed storage for data security," *Proc. of IEEE Inf. Theory Workshop (ITW)*, pp. 506–510, Nov. 2014.

[2] A. Subramanian and S. W. McLaughlin, "MDS codes on the erasure-erasure wiretap channel," *arXiv:0902.3286 [cs.IT]*, 2009.

[3] N. Cai and R. W. Yeung, "Secure network coding on a wiretap network," *IEEE Trans. Inf. Theory*, IT-57, No. 1, pp. 424–435, Jan. 2011.

[4] S. El Rouayheb, E. Soljanin, and A. Sprintson, "Secure network coding for wiretap networks of type II," *IEEE Trans. Inf. Theory*, IT-58, No. 3, pp. 1361–1371, Mar. 2012.

[5] G. R. Blakley, "Safeguarding cryptographic keys," *Proc. of NCC AFIPS*, Vol. 48, pp. 313–317, Jun. 1979.

[6] A. Shamir, "How to share a secret," *Commun. ACM*, Vol. 22, No. 11, pp. 612–613, Nov. 1979.

[7] G. R. Blakley and C. Meadows. "Security of Ramp Schemes," *Advances in Cryptology*, Vol. 196, pp. 242–268, 1985.

[8] N. Merhav and E. Arikan, "The Shannon cipher system with a guessing wiretapper," *IEEE Trans. Inf. Theory*, IT-45, No. 6, pp. 1860–1866, Sep. 1999.

[9] E. Arikan, "An inequality on guessing and its applications to sequential decoding," *IEEE Trans. Inf. Theory*, IT-42, No. 1, pp. 99–105, Jan. 1996.

[10] C. Bunte and A. Lapidoth, "Encoding tasks and Rényi entropy," *IEEE Trans. Inf. Theory*, IT-60, No. 9, pp. 5065–5076, Sep. 2014.

[11] S. Arimoto, "Information measures and capacity of order α for discrete memoryless channels," *Topics in Inf. Theory*, Vol. 17, No. 6, pp. 41-52, 1977.

[12] E. Arikan and N. Merhav, "Guessing subject to distortion," *IEEE Trans. Inf. Theory*, IT-44, No. 3, pp. 1041–1056, May 1998.

[13] C. Bunte and A. Lapidoth, "On the listsize capacity with feedback ," *IEEE Trans. Inf. Theory*, IT-60, No. 11, pp. 6733–6748, Nov. 2014.

[14] C. E. Shannon, Communication theory of secrecy systems, Bell Syst. Tech. J., Vol. 28, No. 3, pp. 565-715, Oct. 1949

[2] The coding scheme is reminiscent of the coding scheme in [1], where after describing X Alice stores part of the description (insecurely) on the first hint, another part (insecurely) on the second hint, and the remaining portion (securely) so that it can only be computed from both hints.