

Related-Key Statistical Cryptanalysis

Darakhshan J. Mir

Department of Computer Science,
Rutgers, The State University of New Jersey

Poorvi L. Vora*

Department of Computer Science,
The George Washington University

Abstract

This paper studies the information-theoretic limits of block cipher statistical key-recovery attacks, which typically use several known plaintext/ciphertext (P/C) pairs to determine a single key. In particular, it studies related-key statistical key recovery, where the adversary uses n related keys, generated from k independent ones. Unlike classical related-key attacks such as differential related-key cryptanalysis, this attack does not exploit a special structural weakness in the cipher or key schedule, but amplifies the weakness exploited in single-key recovery. Using classical results from information theory the paper shows that there exists a relationship among the keys for which the number of P/C pairs required per independent key bit is finite, for any probability of key-recovery error. This may be compared to the unbounded number required per bit of the single-key-recovery attack; the adversarial advantage being similar to that of using error-correcting codes instead of repetition codes for channel communication. The paper also provides lower bounds on the number of P/C pairs required per independent key bit. The practical implications of the results are demonstrated through experiments on reduced-round DES.

Keywords: related keys, concatenated codes, communication channel, statistical cryptanalysis, linear cryptanalysis, DES, channel coding theorem

*Corresponding Author: 801 22nd. St. NW, Washington DC 20052; poorvi@gwu.edu

1 Introduction

Statistical key-recovery attacks (such as linear [16] or differential [3] cryptanalysis) typically require a large number of ciphertexts to successfully estimate the key. Because of this, it is generally assumed that changing the key often offers good protection against such attacks. This is clearly true if the different keys are independent; however, relationships among keys can arise in a number of situations: when the random number generators used in key generation are weak, or when the adversary is powerful enough to control the relationship. While formal models of block cipher cryptanalysis [10, 23, 27] and of related-key attacks [1] exist, there is no model of the combination. In particular, it is not known to what degree the relationship among keys affects the success probability of a statistical attack.

Consider a well-known statistical (single-)key-recovery attack, linear cryptanalysis. Suppose the adversary uses linear cryptanalysis on reduced-round DES to obtain 12 of the 54 keys bits as described by Matsui in [15]. Suppose he uses N P/C pairs to obtain an error of $\epsilon(N)$ in recovering these 12 bits (and hence also the rest of the key, using exhaustive search). Suppose k independent keys are used over a certain period of time; the error in estimating each key is $\epsilon(N)$, and the probability that at least one of the keys is incorrectly estimated is $\varepsilon = 1 - (1 - \epsilon(N))^k$. Further, the communication complexity of the attack, per independent key bit, is $\nu = \frac{N}{d}$. We refer to this attack as an independent-key-recovery attack (defined more formally in section 4), and note that:

$$\varepsilon \rightarrow 0 \Rightarrow N \rightarrow \infty \Rightarrow \nu \rightarrow \infty \tag{1}$$

Now consider the following example of the combination of a related-key attack with linear cryptanalysis. Suppose that $k = 3$, and the adversary uses $n = 7$ related keys, instead of $k = 3$ independent keys. The related keys are generated from the independent keys using the following simple error-correction code that corrects one error. If \check{l}_1, \check{l}_2 and \check{l}_3 are the three independent keys and $\check{k}_1, \check{k}_2, \dots, \check{k}_7$ the related ones,

$$\begin{aligned} \check{k}_i &= \check{l}_i \quad \text{for } i = 1, 2, 3 \\ \check{k}_4 &= \check{l}_1 \oplus \check{l}_2; & \check{k}_5 &= \check{l}_2 \oplus \check{l}_3; & \check{k}_6 &= \check{l}_3 \oplus \check{l}_1; \\ \check{k}_7 &= \check{l}_1 \oplus \check{l}_2 \oplus \check{l}_3 \end{aligned}$$

where \oplus denotes bitwise XOR. To obtain the same value of ν , $\frac{3N}{7}$ P/C pairs are used per related key. Each related key will be determined with a probability of error $\epsilon(\frac{3N}{7}) > \epsilon(N)$. However, a single error in the estimation of the j^{th} bit of all seven keys can be detected and corrected because of the known relationship among the keys, for any value of j .

This paper studies whether the related-key statistical key-recovery attack (defined more formally in section 4) presents any advantages to the adversary, over the independent-key attack. In particular, it studies the behavior of ε to determine if (1) continues to be valid for related-key recovery. It shows that there exists a deterministic relationship among the keys (an error correcting code), for which the power of the single-key recovery is considerably amplified, in a manner made more precise below. It also obtains bounds on adversary efficiency.

1.1 Contributions

The contributions of this paper are threefold:

- It defines the general known-plaintext statistical single-key-recovery attack, and presents a *cryptanalytic channel model (CCM)* for it. The model treats single-key-recovery attacks as

communication over a low capacity channel, using an encoding determined by the cipher and the attack.

- It defines a new attack – the related-key-recovery attack – for all ciphers already vulnerable to single-key recovery. This attack corresponds to a concatenated code in the CCM. It shows the following:
 - For any positive value of ε , however small, there exist values of k (the number of independent keys) and n (the number of related keys) and a relationship among the keys such that $\nu \simeq \Lambda$ for some constant finite Λ . On the other hand, the value of ν is unbounded for independent-key-recovery; the difference is similar to that between using channel codes and repetition codes for channel communication.
 - Λ is asymptotically bounded below. A tight bound is obtained for the case where the n key estimates are made independently of each other, while a loose bound is obtained for the general case. The bounds are similar to the upper bound of channel capacity on the rate of a channel code.
- It provides experimental results for related-key linear cryptanalysis on reduced-round DES using Reed-Solomon codes for the key relationships.
 - An assumption in the related-key attack—that the error in estimating a key is independent of the value of the key, and relationships among keys do not imply strong relationships among the corresponding key estimation errors—has not been previously validated in the literature. Hence, while error expressions for Reed-Solomon codes are available, it was not possible to use them directly to determine the errors of the related-key attacks. However, the experimental results presented in this paper were compared to the Reed-Solomon error expressions and found to be very close; hence these results also validate the assumption.
 - Example results are as follows. The use of a (15, 11) Reed-Solomon code for the key relationship, and $\varepsilon = 0.04$, has a value of ν that is about 11% smaller than that for the corresponding independent-key attack. As another example, for a (127, 87) Reed-Solomon code, the same value of ν provides $\varepsilon = 0.17$ for the related-key attack, and $\varepsilon = 0.81$ for the independent-key attack.

From our results, it follows that, if n and k are large enough and ε is small enough, the value of ν can be made as small a fraction of that of independent-key recovery as desired. Thus the adversary can be at an advantage when the keys are changed more frequently (n times) but are related, than if they are changed less frequently (k times) but are independent. While this is known to be true in several other cases, such as in differential related-key cryptanalysis, our results show that this is true for *any kind* of statistical key-recovery attack, and not simply one that identifies a structural weakness in a cipher or key schedule. While this paper examines the case of the keys being related deterministically, the techniques described here should be useful in various other settings where weaker key relationships are examined, including in the design of key schedules.

The framework of this paper is one of unique key estimates; however, it is fairly common in cryptanalysis to obtain a small list of key estimates, ranked by the value of the likelihood function. In section 4.4 the paper also describes how the related-key-recovery attack, based on unique single-key estimates, can provide an improvement over independent-key attacks that are based on lists of single-key estimates, as long as $N \rightarrow \infty$ as $\varepsilon \rightarrow 0$. A fairer comparison would be one where both

independent-key attacks and related-key attacks use lists of key estimates instead of unique key estimates; however this is outside the scope of this paper.

An early version of this work was presented as an extended abstract in [26], which addressed only the theoretical results, (providing only proof sketches) and only for linear cryptanalysis. This paper generalizes the results to statistical cryptanalysis, provides complete proofs, and provides the results of experimental verification. All the experimental verification is described in detail in [19].

1.2 Organization

This paper is organized as follows. Section 2 presents related work. Section 3 defines the general statistical single-key-recovery attack and the *cryptanalytic channel model (CCM)*. Section 4 describes and defines the related-key-recovery attack, and proves the main theoretical result of the paper – that the related-key-recovery attack can provide a constant value of ν for any value of ε . Section 5 presents experimental results on reduced-round DES. Section 6 presents conclusions and directions for future research.

2 Related Work

The framework of Wagner [27] describes the techniques for obtaining the probabilistic relationships among the plaintext, ciphertext and key for block cipher statistical cryptanalysis. It models the relationships as Markov chains, in the manner of Vaudenay [23, 24].

Stream cipher cryptanalysis was first modeled as channel communication in the late eighties by Meier and Staffelbach [17, 18], for the purpose of constructing efficient attacks. Jakobsen was perhaps the first to explicitly use coding theory in block cipher cryptanalysis, and to treat cryptanalysis of block ciphers as channel communication. Specifically, [11, 9, 10, 12] examined a known plaintext attack on block ciphers where ciphertext can be approximated by a key-dependent polynomial in the plaintext. The fixed-key attack was modeled as channel communication, where the channel communicates the polynomial coefficients, encoded with a Reed-Solomon code. Efficient list decoding was used to obtain estimates of the encryption function.

In Filiol’s model for ciphertext-only attacks [5], the input to the channel is a single binary property of the key. Its output is the parity of a few bits of the ciphertext. The channel output is equal to the channel input with a probability slightly greater than half. Each use of the cipher transmits the same property over the channel, and corresponds to a repetition code on the property. [5] also describes how the same set of N received bits may be decoded as a single repetition code of length N , or as n codes of length $\frac{N}{n}$. This is the decoding technique for a concatenated code, with an inner repetition code of length $\frac{N}{n}$ (over the property of the key), and an outer repetition code of length n (over the key). Filiol correctly indicates that, in this case, concatenation provides no advantage, and that the most efficient decoding is one where the received bits are treated as consisting of a single codeword. While Filiol’s claim of a successful AES break was shown to be incorrect, the approach is very useful to unify block cipher statistical attacks.

Biham examines related-key attacks on block ciphers, tracing the relationships among the keys to the key scheduling algorithm [2]. Kelsey, Schneier and Wagner [13, 14] present related-key attacks on various block ciphers, and demonstrate how real protocols can be exploited to mount such attacks. Bellare and Kohno [1] describe how the most general related-key attack can be very powerful.

No prior single framework addresses both related-key recovery and statistical cryptanalysis. The *CCM* presented in this paper extends the model of Filiol to include known-plaintext attacks and related-key attacks. While Filiol uses concatenation only for decoding, this paper uses it for

the purpose of increasing the efficiency of transmission across the cipher channel. In contrast to the Markov chain models of Vaudenay and of Wagner, the *CCM* models the relationship among plaintext and ciphertext as a communication channel. This allows the *CCM* to address related-key attacks, and also allows access to the coding theory literature. At the same time, the *CCM* allows, in a very natural way, the use of the Markov chain models to determine the communication channel, and the properties transmitted across it.

3 The Known-Plaintext Statistical Single-Key-Recovery Attack

In this section we present our framework. We define the general statistical attack on a block cipher, and present the cryptanalytic channel model. We also show how several common attacks satisfy the definition of the statistical attack, and describe the cryptanalytic channel for these attacks.

Our notation is defined as needed. In general, upper-case letters denote random variables (r.v.s), and lower-case letters specific values taken by the r.v.s. Boldface letters denote sets of r.v.s.

3.1 Definition

We consider known-plaintext statistical key-recovery attacks (such as linear, differential, noisy-polynomial and integral cryptanalysis) on block ciphers. The plaintext and ciphertext are denoted X and Y respectively, and are drawn from the set of q -bit strings, Σ^q . The key is denoted \check{K} , and is drawn from keyspace \mathcal{K} , the set of b -bit strings. The adversary is able to obtain N sets of observations of X and Y , denoted $\{(x_j, y_j)\}_{j=1}^N$, for a fixed key $\check{K} = \check{k}$. These are generated by picking x_j uniformly at random and encrypting it using the block cipher and key \check{k} to obtain y_j . In the case of attacks such as differential and integral cryptanalysis, a single observation consists of more than one P/C pair, and the plaintexts in a single observation are related in a specific manner. In such cases, a single observation consists of a sequence of plaintext values, \mathbf{X} , and an associated sequence of ciphertext values, \mathbf{Y} .

A key-recovery attack requires a random variable S – a function of observable random variables X and Y – whose distribution leaks information about the key. In general, the most probable value of S is random variable T , a function of one or more bits of \check{k} . The adversary uses $\{(x_j, y_j)\}_{j=1}^N$ to obtain the value of S , and, through this, a maximum likelihood estimate of the key bits, assuming, as is typical in estimation theory (see, for example, [22]), a uniform *a priori* distribution on \check{K} .

The general known-plaintext statistical single-key-recovery attack on block ciphers may be defined as follows:

Definition 1 *A known-plaintext statistical single-key-recovery attack on a block cipher with plaintext X and ciphertext Y encrypted with fixed key $\check{k} \in \mathcal{K}$ consists of:*

- *Function $\kappa, \kappa : \mathcal{K} \rightarrow \kappa(\mathcal{K}) \subseteq \mathcal{K}$*
- *A function Many mapping a single plaintext X to a sequence of plaintexts \mathbf{X} , $\text{Many}(X) = \mathbf{X}$. The corresponding sequence of ciphertexts is denoted \mathbf{Y}*
- *Random variables $S(\mathbf{X}, \mathbf{Y}) \in \mathcal{Z}$ and $T(\mathbf{X}, \mathbf{Y}, \kappa(\check{k})) \in \mathcal{Z}$, where \mathcal{Z} denotes a domain of size m*
- *N instances of (\mathbf{X}, \mathbf{Y}) : $\{(\mathbf{x}_j, \mathbf{y}_j)\}_{j=1}^N$*
- *Algorithm KeyRecovery*

such that:

- $|\kappa(\mathcal{K})| = 2^d$, where $|\cdot|$ denotes size
- $S(\mathbf{X}, \mathbf{Y})$ takes on the value $T(\mathbf{X}, \mathbf{Y}, \kappa(\ddot{k}))$ (slightly) more often than it takes on any other value in \mathcal{Z} when X is uniformly distributed. Further, it takes on all other values with equal probability. Hence, for $Z \in \mathcal{Z}$:

$$\Pr[S(\mathbf{X}, \mathbf{Y}) = Z] = \begin{cases} \frac{1}{m} + \gamma & Z = T(\mathbf{X}, \mathbf{Y}, \kappa(\ddot{k})) \\ \frac{1}{m} - \frac{\gamma}{m-1} & \text{else} \end{cases}$$

for small positive γ .

- Algorithm *KeyRecovery* provides estimate(s) $\widehat{\kappa(\ddot{k})}$ of $\kappa(\ddot{k})$:

$$\widehat{\kappa(\ddot{k})} = \underset{z \in \kappa(\mathcal{K})}{\text{max}} \mid \{(\mathbf{x}_j, \mathbf{y}_j) : S(\mathbf{x}_j, \mathbf{y}_j) = T(\mathbf{x}_j, \mathbf{y}_j, z)\} \mid$$

and is considerably more efficient than an exhaustive search over all possible values of \ddot{k} .

The attack is denoted $\Gamma = (\kappa, d, \text{Many}, S, T, N, \gamma, \text{KeyRecovery})$.

Note that S may play the part of a statistical distinguisher (a random variable that is uniformly distributed when X and Y are independent of each other and uniformly distributed, but is non-uniformly distributed when Y is obtained by encrypting uniformly distributed X using the block cipher). However, it is not sufficient for S to be a statistical distinguisher. For the purpose of key-recovery, the distribution of S should reveal information on \ddot{k} .

Note also that we do not specify what is meant by algorithm *KeyRecovery* being “considerably more efficient than an exhaustive search”. Our condition on the algorithm is required to eliminate trivial and exorbitantly expensive attacks – an example is when $S = Y$, $T = E_{\ddot{k}}(X)$, $\gamma = \frac{n-1}{n}$, and \ddot{k} is determined by an exhaustive search over \mathcal{K} . The efficiency requirement for *KeyRecovery* may be met by the use of a compression κ – as in example 1 – such that an exhaustive search over all possible values of $\kappa(\ddot{k})$ is not prohibitive. It may also be met by the use of r.v.s T and S such that a maximum-likelihood estimate of $\kappa(\ddot{k})$ may be obtained more efficiently than through an exhaustive search of \mathcal{K} .

Finally, we define the key-recovery error and the amortized communication cost of the attack.

Definition 2 The probability of key-recovery error of the known-plaintext statistical single-key-recovery attack Γ is:

$$\epsilon(N) = \Pr[\kappa(\ddot{k}) \neq \widehat{\kappa(\ddot{k})}]$$

Definition 3 The amortized communication cost, ν , of a statistical single-key-recovery attack in P/C pairs used per bit is:

$$\nu = \frac{N}{d}$$

3.2 The Cryptanalytic Channel Model

The key-recovery attack of Definition 1 may be described as channel communication using the cryptanalytic channel model (CCM). First, we define the *cryptanalytic channel* over which the communication occurs.

Definition 4 The cryptanalytic channel of the single-key-recovery attack Γ is the simplex communication channel with probability of error $p_e = \frac{n-1}{n} - \gamma$, and input and output alphabets \mathcal{Z} .

Recall that a simplex channel is one in which the probability of error does not depend on the input or output symbols.

$\kappa(\ddot{k})$ forms the message (see Figure 1). The adversary can observe $S(\mathbf{X}, \mathbf{Y})$, which forms the channel output and is a noisy value of a P/C-dependent key property, $T(\mathbf{X}, \mathbf{Y}, \kappa(\ddot{k}))$. The randomness is provided by the value of P , chosen uniformly at random. Thus $[T(\mathbf{X}_1, \mathbf{Y}_1, \kappa(\ddot{k})), T(\mathbf{X}_2, \mathbf{Y}_2, \kappa(\ddot{k})), \dots, T(\mathbf{X}_N, \mathbf{Y}_N, \kappa(\ddot{k}))]$ forms the transmitted codeword, of length N . $[S(\mathbf{X}_1, \mathbf{Y}_1), S(\mathbf{X}_2, \mathbf{Y}_2), \dots, S(\mathbf{X}_N, \mathbf{Y}_N)]$ forms the received codeword. The decoding algorithm is Algorithm *KeyRecovery*, and corresponds to maximum-likelihood decoding. The rate of communication, R , is the inverse of the communication complexity per key bit, which is $\frac{N}{d}$.

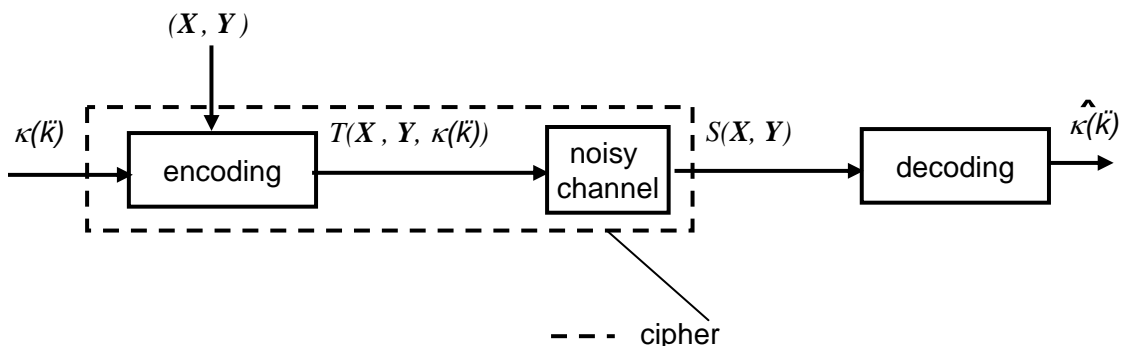


Figure 1: The Cryptanalytic Channel Model

Definition 1 does not specifically assume an iterated cipher. In the following we focus on a specific common type of key recovery attack in an r -round iterated block cipher: the determination of the r^{th} round key, $\ddot{k}^{(r)}$, using a statistical distinguisher for $r - 1$ rounds.

3.3 $\ddot{k}^{(r)}$ -recovery

We now consider an r -round iterated block cipher. The message input to the i^{th} round is denote $X^{(i)}$, its output $Y^{(i)}$, and the i^{th} round key is denoted $\ddot{k}^{(i)}$. An attack that determines $\ddot{k}^{(r)}$ is a special kind of statistical key recovery attack. It is based on the existence of a random variable A that is a statistical distinguisher for the $(r - 1)$ -round cipher [8]. That is, A is a function of X and $Y^{(r-1)}$, uniformly distributed when X and $Y^{(r-1)}$ are independent of each other and uniformly distributed, but non-uniformly distributed when $Y^{(r-1)}$ is obtained by encrypting uniformly distributed X using a fixed key. The most likely value of A , denoted B , may be a function of X , Y and/or the key for the first $r - 1$ rounds, denoted \ddot{k}_{-r} . The bit-length of \ddot{k}_{-r} is denoted b_r . Hence:

$$Pr[A(\mathbf{X}, \mathbf{Y}^{(r-1)}) = Z] = \begin{cases} \frac{1}{m} + \gamma & Z = B(\mathbf{X}, \mathbf{Y}^{(r-1)}, \kappa_{-r}(\ddot{k}_{-r})) \\ \frac{1}{m} - \frac{\gamma}{m-1} & \text{else} \end{cases}$$

for $\gamma > 0$, and some function κ_{-r} . As A is a distinguisher, B can be constant as a function of \ddot{k}_{-r} . Observing that $Y^{(r-1)} = \mathcal{F}_{\ddot{k}^{(r)}}^{-1}(Y)$, where $\mathcal{F}_{\ddot{k}^{(i)}}$ is the round function with round key $\ddot{k}^{(i)}$, and rearranging to obtain an observable random variable (a function of only \mathbf{X} and \mathbf{Y}) on the left and

a function of \mathbf{X} , \mathbf{Y} , \ddot{k}_{-r} and $\ddot{k}^{(r)}$ on the right, we get:

$$Pr[S'(\mathbf{X}) = Z] = \begin{cases} \frac{1}{m} + \gamma & Z = T'(\mathbf{X}, \mathcal{F}_{\ddot{k}^{(r)}}^{-1}(\mathbf{Y}), \kappa_{-r}(\ddot{k}_{-r})) \\ \frac{1}{m} - \frac{\gamma}{m-1} & \text{else} \end{cases}$$

for some functions S' and T' if we may assume that error probabilities do not depend on the input symbol after the rearrangement of terms. Note that, if T' is constant as a function of \ddot{k}_{-r} (such as in integral and differential cryptanalysis, see section 3.4), κ_{-r} is a trivial function. Note also that S' may be constant as a function of \mathbf{X} .

$\hat{\ddot{k}}^{(r)}$ and $\widehat{\kappa_{-r}(\ddot{k}_{-r})}$ are the maximum-likelihood estimates:

$$\widehat{\kappa_{-r}(\ddot{k}_{-r})}, \hat{\ddot{k}}^{(r)} = \underset{w, z}{\operatorname{argmax}} |\{(\mathbf{x}_i, \mathbf{y}_i) : S'(\mathbf{X}) = T'(\mathbf{X}, \mathcal{F}_z^{-1}(\mathbf{Y}), w)\}|$$

It is thus clear that a $\ddot{k}^{(r)}$ -recovery attack satisfies Definition 1 of a statistical key recovery attack.

Viewing $\ddot{k}^{(r)}$ -recovery in the CCM, we observe the following. The message is $\{\kappa_{-r}(\ddot{k}_{-r}), \ddot{k}^{(r)}\}$ or simply $\ddot{k}^{(r)}$ (if T' is independent of \ddot{k}_{-r}), and the j^{th} code symbol is denoted \mathcal{I}_j :

$$\mathcal{I}_j(\kappa_{-r}(\ddot{k}_{-r}), \ddot{k}^{(r)}) = T'(\mathbf{x}_j, \mathcal{F}_{\ddot{k}^{(r)}}^{-1}(\mathbf{y}_j), \kappa_{-r}(\ddot{k}_{-r}))$$

The codeword, α , is: $\alpha = [\mathcal{I}_1(\kappa_{-r}(\ddot{k}_{-r}), \ddot{k}^{(r)}), \mathcal{I}_2(\kappa_{-r}(\ddot{k}_{-r}), \ddot{k}^{(r)}), \dots, \mathcal{I}_N(\kappa_{-r}(\ddot{k}_{-r}), \ddot{k}^{(r)})]$ The received codeword is $[S'(\mathbf{x}_1), S'(\mathbf{x}_2), \dots, S'(\mathbf{x}_N)]$

3.4 Examples: Linear and Differential Cryptanalysis

In this section we illustrate the model with the examples of linear and differential cryptanalysis.

3.4.1 Linear Cryptanalysis

Consider Matsui's $k^{(r)}$ -recovery attack based on linear cryptanalysis ([15], "Algorithm 2"). The attack is based on the following statistical distinguisher, A :

$$A(X, Y^{(r-1)}) = f_{r-1}(X) \oplus g_{r-1}(Y^{(r-1)})$$

and $B(\mathbf{X}, \mathbf{Y}^{(r-1)}, \kappa_{-r}(\ddot{k}_{-r})) = h_{r-1}(\ddot{k}_{-r})$ for linear/affine functions f_{r-1} , g_{r-1} , h_{r-1} . The following probabilistic relationship is hence known:

$$Pr[f_{r-1}(X) = g_{r-1}(\mathcal{F}_{\ddot{k}^{(r)}}^{-1}(Y)) \oplus h_{r-1}(\ddot{k}_{-r})] = \frac{1}{2} + \gamma \quad (2)$$

for some $\gamma > 0$.

$\widehat{\ddot{k}}^{(r)}$ is the value z that satisfies $f_{r-1}(x_j) = g_{r-1}(\mathcal{F}_z^{-1}(y_j))$ most or least often. That is, if $\phi(z) = |\{(x_j, y_j) : f_{r-1}(x_j) = g_{r-1}(\mathcal{F}_z^{-1}(y_j))\}|$,

$$\widehat{\ddot{k}}^{(r)} = \underset{z}{\operatorname{argmax}} \left\| \phi(z) - \frac{N}{2} \right\|$$

where $\|\cdot\|$ denotes the absolute value. If $\phi(\widehat{\ddot{k}}^{(r)}) > \frac{N}{2}$, $h_{r-1}(\widehat{\ddot{k}}^{(r)}) = 0$, else $h_{r-1}(\widehat{\ddot{k}}^{(r)}) = 1$. This simple algorithm is equivalent to maximum likelihood estimation [7]. A list of estimates of $\ddot{k}^{(r)}$ and $h_{r-1}(\ddot{k}_{-r})$ may also be obtained, ranked in order of the corresponding values of $\|\phi(z) - \frac{N}{2}\|$.

Hence $(r - 1)$ -round linear cryptanalysis satisfies Definition 1 with

$$\begin{aligned} \mathcal{Z} &= \mathbb{Z}_2; & \mathbf{X} &= X; & \kappa(\ddot{k}) &= h_{r-1}(\ddot{k}_{-r}) \parallel \ddot{k}^{(r)}; & d &= b_r + 1; \\ S(X, Y) &= f_{r-1}(X); & T(X, Y, \ddot{k}) &= g_{r-1}(\mathcal{F}_{\ddot{k}^{(r)}}^{-1}(Y)) \oplus h_{r-1}(\ddot{k}_{-r}) \end{aligned}$$

Further, it is known that [16, Lemmas 2 and 5]:

$$\epsilon \rightarrow 0 \Rightarrow N \rightarrow \infty \tag{3}$$

In the CCM, the r^{th} round key, $\ddot{k}^{(r)}$, and one bit of the rest of the key, $h_{r-1}(\ddot{k}_{-r})$, form the message. The transmitted codeword is of size N , where the j^{th} bit, denoted $\mathcal{I}_j(\ddot{k}^{(r)}, h_{r-1}(\ddot{k}_{-r}))$ is:

$$\mathcal{I}_j(\ddot{k}^{(r)}, h_{r-1}(\ddot{k}_{-r})) = g_{r-1}(\mathcal{F}_{\ddot{k}^{(r)}}^{-1}(y_j)) \oplus h_{r-1}(\ddot{k}_{-r})$$

(see (2)). Notice that the codeword is ciphertext dependent. Notice also that it is non-linear in the bits of $\ddot{k}^{(r)}$, though linear/affine in the bits of \ddot{k}_{-r} .

The codeword itself is not accessible to the adversary. However, $[(f_{r-1}(x_1), f_{r-1}(x_2), \dots, f_{r-1}(x_N))]$ is a very noisy value of the codeword, providing the output of the cryptanalytic channel (see (2)). The channel error probability is $p_e = \frac{1}{2} - \gamma$, the channel is symmetric, and the corresponding capacity is $\mathcal{C} \simeq \frac{2\gamma^2}{\ln 2}$ (the first non-zero term in the Taylor series expansion). $\ddot{k}^{(r)}$ and $h_{r-1}(\ddot{k}_{-r})$ are determined from the values of $f_{r-1}(x_j)$ using maximum-likelihood decoding [20].

3.4.2 Differential Cryptanalysis

In the differential cryptanalytic attack, there exist values ΔX and ΔY such that, after $r - 1$ rounds of the cipher, a difference of ΔX in plaintext results in a difference ΔY in $Y^{(r-1)}$, more often than in the ideal random cipher. If two plaintext values X_1 and X_2 are encrypted with key \ddot{k} to give ciphertext Y_1 and Y_2 , and $X_1 \oplus X_2 = \Delta X$,

$$Pr[\Delta Y = \mathcal{F}_{\ddot{k}^{(r)}}^{-1}(Y_1) \oplus \mathcal{F}_{\ddot{k}^{(r)}}^{-1}(Y_2)] = \frac{1}{2^q} + \gamma \tag{4}$$

for some $\gamma > 0$, as $Y_i^{(r-1)} = \mathcal{F}_{\ddot{k}^{(r)}}^{-1}(Y_i)$, $i = 1, 2$. Differential cryptanalysis satisfies Definition 1 with $\mathcal{Z} = \Sigma^q$, $\mathbf{X} = \{X, X \oplus \Delta P\}$, $\kappa(\ddot{k}) = \ddot{k}^{(r)}$, $d = b_r$, $S(\mathbf{X}, \mathbf{Y}) = \Delta Y$, $T(\mathbf{X}, \mathbf{Y}, \kappa(\ddot{k})) = \mathcal{F}_{\ddot{k}^{(r)}}^{-1}(Y_1) \oplus \mathcal{F}_{\ddot{k}^{(r)}}^{-1}(Y_2)$.

Other common types of block cipher statistical cryptanalysis—such as noisy polynomial, differential and integral cryptanalysis—may also be represented as channel communication.

4 Related-Key Recovery

As statistical cryptanalysis corresponds to channel communication, the communication efficiency of an attack may be studied using classical results from information theory. In particular, the channel coding theorem provides a tight bound on the adversary's efficiency. In communication theory, the repetition code, which consists of the transmission of a single symbol over the channel n times, requires a decrease in rate (number of information symbols per code symbol transmitted) for a decrease in error. Note that this is similar to the fact that a single-key statistical attack also requires an increase in communication cost (P/C pairs) for a decrease in error. On the other hand, in communication theory, the transmission of related symbols can provide a very different relationship between rate and error. Shannon's channel coding theorem [21, 4] proves that there

exists a channel code, using which the probability of decoding error can be made as small as desired, for constant rate R , as long as n and k can be increased indefinitely, and $R \leq \mathcal{C}$. Further, it says that this is not possible for $R > \mathcal{C}$.

Applied to the cryptanalytic channel, the channel coding theorem says that ν can be maintained at a constant value, while ϵ is made as small as desired, as long as N and d can be increased indefinitely, and ν – which is the inverse of R – is at least as great as the inverse of \mathcal{C} . As d is constant and determined by the single-key recovery attack, it is not possible to increase d indefinitely. Further, to achieve efficient transmission, T would have to correspond to the Shannon code that achieves the limits of the channel coding theorem, and T too is determined by the attack. Thus the channel coding theorem cannot directly be applied to the problem of communicating across the cryptanalytic channel. It is however, possible to apply the channel coding theorem to the *superchannel* – one may consider $\kappa(\ddot{k})$ of single-key recovery as being transmitted to the adversary through the cryptanalytic attack, which may be thought of as a superchannel, with probability of error ϵ . The adversary receives $\widehat{\kappa(\ddot{k})}$, transmitted once for one single-key recovery attack.

Definition 5 *The superchannel of known-plaintext statistical single-key-recovery attack Γ is a simplex channel with input and output alphabet $\kappa(\mathcal{K})$ and probability of error $\epsilon(N)$.*

To communicate efficiently across the superchannel, the adversary would encode the values of $\kappa(K)$; that is, the adversary would encode independent keys, which form the message, to obtain channel-coded keys (while the channel-coding needs to be performed on the values of $\kappa(\ddot{k}_i)$, the same code may be used for the entire key, with the only caveat being that the bits representing $\kappa(\ddot{k}_i)$ must be encoded separately from the other bits). The channel-coded keys would then be transmitted across the superchannel – that is, each of the channel-coded keys would be used in a single-key recovery attack. This intentionally assumes a very powerful adversary. The results of the channel coding theorem provide upper bounds on the efficiency of even this very powerful adversary, and the channel coding model is a useful one for examining the effect of key relationships on the value of ν . It is well-known in coding theory that efficient communication over a noisy channel is not obtained by the communication of independent message symbols. Thus, it is natural that a relationship among keys will provide a decrease in ν .

In this section we describe a general related-key attack. We prove that the use of related keys can reduce considerably the amortized cost of the attack.

4.1 The General Statistical Related-Key-Recovery Attack

Definition 6 *A known-plaintext statistical related-key-recovery attack on a block cipher with plaintext X , ciphertext Y , keyspace \mathcal{K} and key $\ddot{k} \in \mathcal{K}$ consists of:*

- A statistical single-key-recovery attack Γ
- k independent keys, $\ddot{\mathbf{l}} = (\ddot{l}_1, \ddot{l}_2, \dots, \ddot{l}_k)$, such that $\ddot{l}_i \in \mathcal{K} \forall i$
- n related keys, $\ddot{\mathbf{k}} = (\ddot{k}_1, \ddot{k}_2, \dots, \ddot{k}_n)$, such that $\ddot{k}_i \in \mathcal{K} \forall i$ and $n \geq k$
- An injective function ψ (the encoding of $\kappa(\ddot{k})$), such that:

$$\begin{aligned} \psi : (\kappa(\mathcal{K}))^k &\rightarrow (\kappa(\mathcal{K}))^n \\ \psi(\kappa(\ddot{\mathbf{l}})) &= \kappa(\ddot{\mathbf{k}}) \end{aligned}$$

where $\kappa(\ddot{\mathbf{l}})$ and $\kappa(\ddot{\mathbf{k}})$ denote $(\kappa(\ddot{l}_1), \kappa(\ddot{l}_2), \dots, \kappa(\ddot{l}_k))$ and $(\kappa(\ddot{k}_1), \kappa(\ddot{k}_2), \dots, \kappa(\ddot{k}_n))$ respectively

- Algorithm `RelatedKeyRecovery` that obtains estimates $\widehat{\kappa(\check{k}_i)}$ independently $\forall i$, using the single key recovery attack, and uses $\widehat{\kappa(\check{\mathbf{k}})} = (\widehat{\kappa(\check{k}_1)}, \widehat{\kappa(\check{k}_2)}, \dots, \widehat{\kappa(\check{k}_n)})$ to obtain maximum likelihood estimate $\widehat{\kappa(\check{\mathbf{l}})} = (\widehat{\kappa(\check{l}_1)}, \widehat{\kappa(\check{l}_2)}, \dots, \widehat{\kappa(\check{l}_k)})$.

A statistical related-key-recovery attack is denoted $\Psi = (\Gamma, k, n, \psi, \text{RelatedKeyRecovery})$. As we will be comparing the use of related keys to the use of independent keys, we now define the independent-key attack.

Definition 7 A known-plaintext independent-key-recovery attack is a known-plaintext related-key-recovery attack with $k = n$ and $\psi = I$, the identity.

Finally, we define the amortized communication cost of the related-key attack, and its error.

Definition 8 The amortized communication cost, ν , of a statistical related-key-recovery attack Ψ in P/C pairs used per bit is:

$$\nu(N, d, k, n) = \frac{nN}{kd}$$

Definition 9 The related-key-recovery error, ε , of a statistical related-key-recovery attack Ψ , is the probability that $\widehat{\kappa(\check{\mathbf{l}})} \neq \kappa(\check{\mathbf{l}})$.

$$\varepsilon(N, k, n, \psi) = Pr[\widehat{\kappa(\check{\mathbf{l}})} \neq \kappa(\check{\mathbf{l}})]$$

We observe that, for a given single-key-recovery attack Γ with key-recovery error $\epsilon(N)$ and the corresponding independent-key-recovery attack, $(\Gamma, k, k, I, \text{RelatedKeyRecovery})$, with related-key-recovery error $\varepsilon(N, k, k, I)$, $1 - \varepsilon(N, k, k, I) = (1 - \epsilon(N))^k$.

4.2 Related-Key-Recovery Attacks as Concatenated Codes

In communication theory, the combination of an *inner code*, whose encoding and decoding form part of the superchannel, and an *outer code*, used to transmit over the superchannel, form a *concatenated code*. Consider a code taking k_1 message symbols from alphabet \mathcal{X} to a codeword of length n_1 over \mathcal{X} . This is the inner code, used to transmit over the channel. One may view the entire coding/transmitting/decoding process as a superchannel, over which a single message symbol from \mathcal{X}^{k_1} is transmitted with error equal to the decoding error of the code. One may further encode for efficient transmission over the superchannel, and a message of k_2 symbols, each from \mathcal{X}^{k_1} , may be encoded to a codeword of size n_2 symbols, each also from \mathcal{X}^{k_1} .

Definition 10 A concatenated code is a code h , $h : \mathcal{X}^{k_{conc}} \rightarrow \mathcal{X}^{n_{conc}}$ such that \exists

- f , the inner code, $f : \mathcal{X}^{k_{1,conc}} \rightarrow \mathcal{X}^{n_{1,conc}}$
- g , the outer code, $g : (\mathcal{X}^{k_{1,conc}})^{k_{2,conc}} \rightarrow (\mathcal{X}^{k_{1,conc}})^{n_{2,conc}}$

such that $h = f \circ g$, $k_{conc} = k_{1,conc}k_{2,conc}$, and $n_{conc} = n_{1,conc}n_{2,conc}$.

The decoding algorithm of the concatenated code is the composition of the decoding algorithms of g and f . That is, a received string $m_1m_2m_3\dots m_{n_{conc}} \in \mathcal{X}^{n_{conc}}$ is divided into $n_{2,conc}$ substrings of size $n_{1,conc}$. Each substring is decoded using the decoding algorithm of f to obtain a string of size $k_{1,conc}$. $n_{2,conc}$ strings, each of size $k_{1,conc}$, provide $n_{2,conc}$ symbols from $\mathcal{X}^{k_{1,conc}}$, and are decoded to obtain $k_{2,conc}$ symbols from $\mathcal{X}^{k_{1,conc}}$.

We now have the following simple result.

Lemma 11 *The related-key-recovery attack of Definition 6, Ψ , is a concatenated code over the superchannel of attack Γ .*

Proof. ψ is the outer code; $n_{2,conc} = n$ and $k_{2,conc} = k$. A message encoded by the outer code is of the form $\kappa(\mathbf{\ddot{I}})$ and a single codeword from the outer code is of the form $\kappa(\mathbf{\ddot{k}})$, which forms the input to the superchannel. T with maximum-likelihood decoding forms the inner code, and $n_{1,conc} = N$ and $k_{1,conc} = d$. The output of the superchannel consists of the estimates of the single key-recovery, $\widehat{\kappa(\mathbf{\ddot{I}})}$. The procedure of finding $\widehat{\kappa(\mathbf{\ddot{I}})}$ from $\widehat{\kappa(\mathbf{\ddot{k}})}$ is the maximum likelihood decoding of ψ . ■

4.3 The Existence of an Efficient Related-Key-Recovery Attack

In this section, we prove our main result: that ε can be made as small as desired, while maintaining ν at a constant value. In coding theory terms, this is equivalent to the result that communication error can be made as small as desired, while maintaining the rate of the code at a constant value¹. We approach the problem as follows. In a related-key-recovery attack, we may treat the superchannel with a fixed value of N as a channel over which the outer code is used for communication. While the number of message symbols of the inner code is fixed at d , the outer code is not limited in number of message/code symbols. Hence, by the channel coding theorem [21] the superchannel can be used to communicate with any error, at any rate smaller than its capacity (which depends on N and d), as long as n and k can be large enough. Note, however, that we do not show that the rate can be as large as the inner channel capacity.

More formally: consider any error, $\varepsilon(N)$, reasonably small, in a single-key-recovery attack using N P/C pairs. Assuming that the superchannel is symmetric, let its capacity be $\mathcal{C}_S(N)$. The superchannel may be used to communicate at any fixed rate smaller than $\mathcal{C}_S(N)$ with error as low as desired. This gives us:

Theorem 12 *Consider a family of related-key-recovery attacks $\{\Psi_i\}_{i \in \mathbb{Z}^+}$ with fixed d and N such that $\frac{N}{d} \geq \frac{1}{\mathcal{C}}$, where \mathcal{C} is the capacity of the cryptanalytic channel of the single-key recovery attack Γ , and n_i increasing with i . Denote the capacity of the superchannel by $\mathcal{C}_S(N)$. $\exists \{\Psi_i\}_{i \in \mathbb{Z}^+}$ such that*

$$\lim_{i \rightarrow \infty} \varepsilon(N, k_i, n_i, \psi_i) = 0 \text{ and } \nu(d, N, k_i, n_i) = \Lambda, \forall \Lambda \geq \frac{N}{d\mathcal{C}_S(N)}$$

Proof. The result follows from the application of the channel coding theorem to the outer code, ψ , with error $\varepsilon(N, k, n, \psi)$, and the superchannel with capacity $\mathcal{C}_S(N)$. The channel coding theorem [21, 4] says that \exists a family of outer codes $\{\psi_i\}_{i \in \mathbb{Z}^+}$ such that

$$\lim_{i \rightarrow \infty} \varepsilon(N, k_i, n_i, \psi_i) = 0$$

for all constant $R \leq \mathcal{C}_S(N)$, where $R = \frac{k_i}{n_i}$. As

$$\nu(d, N, k_i, n_i) = \frac{n_i N}{k_i d}$$

¹Note that the setting is different from that of the classical result on concatenated codes by Forney [6], which shows that communication error can be made as small as desired, while maintaining the rate of the code at any constant value smaller than inner channel capacity, if the inner code is a good one. In related-key-recovery, the number of message bits for the inner code, d , is small and fixed, and hence the inner code may not be considered a good one. Hence, the results of [6] may not be directly applied. Additionally, [6] shows that a concatenated code can be used to attain the limits of the channel coding theorem if the inner code can be used to do so, and if the message length of the inner code can be increased indefinitely. This is not true for the cryptanalytic channel because d cannot be increased indefinitely. In fact our experimental results indicate that the maximum rate of transmission is almost half the capacity of the inner channel for $(r - 1)$ -round linear cryptanalysis on 8-round DES.

this implies that

$$\nu(d, N, k_i, n_i) = \frac{N}{dR} = \Lambda, \forall \Lambda \geq \frac{N}{d\mathcal{C}_S(\epsilon)}$$

■

Theorem 12 says that, while ϵ is brought indefinitely close to zero, ν can be maintained at a constant value for related-key recovery. On the other hand, ν is unbounded for independent-key recovery (see (1)). Thus, related-key recovery can reduce the value of ν to as small a fraction of the original value as desired, if error is brought indefinitely close to zero, and n and k made as large as desired. We state this more formally in the following corollaries.

Corollary 13 *Given any $\alpha \in (0, 1)$, and a single-key recovery attack Γ , $\exists N_\alpha$, a corresponding related-key-recovery attack Ψ , and M_α , such that*

$$\varepsilon(M_\alpha, k, n, \psi) = \epsilon(N_\alpha)$$

and

$$\frac{\nu(d, M_\alpha, k, n)}{\nu(d, N_\alpha, k, k)} < \alpha$$

Proof. Choose any $M_\alpha \geq \frac{d}{\epsilon}$ and the corresponding ψ of Theorem 12. M_α will be used for the number of P/C pairs in each single-key recovery for the related-key recovery. ψ will define the relationship among the keys. Let $\mathcal{C}_S(M_\alpha)$ be the capacity of the cryptanalytic channel. Choose some $\Lambda \geq \frac{N}{d\mathcal{C}_S(M_\alpha)}$. This will define the rate of the outer code, or the value of $\frac{k}{n}$ for the related keys. Given α , choose $N_\alpha > \frac{\Lambda d}{\alpha}$. These will be the number of P/C pairs used in single-key recovery for the independent-key recovery. Then, by Theorem 12, $\exists n, k$ such that

$$\varepsilon(M_\alpha, k, n, \psi) = \epsilon(N_\alpha)$$

and

$$\nu(d, M_\alpha, k, n) = \Lambda$$

Further,

$$\frac{\nu(d, M_\alpha, k, n)}{\nu(d, N_\alpha, k, k)} = \frac{\Lambda}{\frac{N_\alpha}{d}} = \frac{\Lambda d}{N_\alpha} < \alpha$$

■

Corollary 14 *Given any $\alpha \in (0, 1)$, and a single-key recovery attack Γ , $\exists N_\alpha$, a related-key-recovery attack Ψ , and M_α , such that*

$$\varepsilon(M_\alpha, k, n, \psi) < \varepsilon(N_\alpha, k, k, I)$$

and

$$\frac{\nu(d, M_\alpha, k, n)}{\nu(d, N_\alpha, k, k)} < \alpha$$

Proof. This follows from Corollary 13,

$$\varepsilon(M_\alpha, k, n, \psi) = \epsilon(N_\alpha) = 1 - (1 - \varepsilon(N_\alpha, k, k, I))^{\frac{1}{k}} < \varepsilon(N_\alpha, k, k, I)$$

■

In Theorem 12 and Corollaries 13 and 14 we have illustrated the increase in capability of the adversary provided by related-key recovery. We now turn to the limits on this capability when the key estimates are unique. The channel coding theorem applied to the superchannel (Corollary 15) and to the cryptanalytic channel (Corollary 16) provides lower bounds on the value of ν if it is held constant and error required to be arbitrarily close to zero; the bound of Corollary 15 is tight, but depends on the specific attack and the cipher. More specifically, it depends on the function $\epsilon(N)$. For linear cryptanalysis on 8-round DES, our experimental results (Section 5) indicate this tight bound is almost twice the bound implied by the cryptanalytic channel capacity (the bound of Corollary 16).

Corollary 15 *For a family of related-key-recovery attacks $\{\Psi_i\}_{i \in \mathbb{Z}^+}$ such that n_i is increasing with i , $\lim_{i \rightarrow \infty} \epsilon(N, k_i, n_i, \psi_i) = 0$, and ν is constant, the minimum value of ν is $\frac{\min N}{d\mathcal{C}_S(N)}$*

Proof. Suppose there is a value of ν that is smaller, and that it corresponds to the use of N_0 P/C pairs in each of the single-key-recovery attacks. Then, in particular, $\nu = \frac{n_i N_0}{k_i d} < \frac{N_0}{d\mathcal{C}_S(N_0)}$ and $\frac{n_i}{k_i} < \frac{1}{\mathcal{C}_S(N_0)}$, or the rate of the outer code, $\frac{k_i}{n_i}$, is larger than the capacity of the superchannel, $\mathcal{C}_S(N_0)$, contradicting the channel coding theorem. ■

Corollary 16 *For a family of related-key-recovery attacks $\{\Psi_i\}_{i \in \mathbb{Z}^+}$ with n_i increasing with i , $\nu(d, N, k_i, n_i) = \text{constant}$, and*

$$\lim_{i \rightarrow \infty} \epsilon(N, k_i, n_i, \psi_i) = 0$$

$$\frac{\min N}{d\mathcal{C}_S(N)} \geq \frac{1}{\mathcal{C}}$$

where \mathcal{C} is the capacity of the cryptanalytic channel corresponding to single-key attack Γ .

Proof. If not, $\{\Psi_i\}_{i \in \mathbb{Z}^+}$ would be an example of constant-rate communication with zero asymptotic error over the cryptanalytic channel at a rate greater than its capacity, \mathcal{C} . This would contradict the channel coding theorem. ■

For the same reasons, the above is also a bound on the value of ν if the n_i related keys are not determined independently in single-key recovery attacks.

Finally, the adversary need not maintain a constant value of ν for the attack. Using a result from [25], it can be shown that the tight bound of Corollary 15 is a tight asymptotic bound for attacks where ϵ is arbitrarily small, but ν is not required to be constant.

Corollary 17 *For a family of related-key-recovery attacks $\{\Psi_i\}_{i \in \mathbb{Z}^+}$ such that n_i increasing with i and $\lim_{i \rightarrow \infty} \epsilon(N, k_i, n_i, \psi_i) = 0$, $\lim_{i \rightarrow \infty} \nu(N, d, k_i, n_i) \geq \frac{N}{d\mathcal{C}_S(N)}$*

Proof. Follows from the fact that [25] $\lim_{i \rightarrow \infty} \epsilon(N, k_i, n_i, \psi_i) = 0 \Rightarrow \lim_{i \rightarrow \infty} \frac{k_i}{n_i} \leq \mathcal{C}$. ■

4.4 Non-unique Key Estimates

The framework of this paper focuses on unique key estimates. Most of the ideas may, however, be applied to the situation common to cryptanalysis: instead of a unique estimate, the adversary obtains a (small) list of estimates of the key, ranked by the value of the likelihood function. In this situation, the adversary performs list decoding, instead of unique decoding, at the output of the channel. ϵ , the probability of error of the single-key attack, is the probability that the correct key does not belong to the list of estimates. If the list is small enough, the value of ν still increases

without bound as error decreases without bound. The independent-key attack also results in a list for each of the k independent keys. $\varepsilon(N, k, k, I)$, the probability of error of the independent-key estimate, is the probability that any of the independent keys does not belong to the corresponding list of estimates, and (1) also holds.

The related-key attack uses the lists of estimates of each of the \tilde{l}_i , along with the value of the likelihood function, to rank combinations of estimates from the lists. This provides a list of estimates of $\tilde{\mathbf{I}}$, and Lemma 11 also holds. However, the number of combinations of the estimates grows prohibitively with the value of k . The problem of comparing the performance of related-key and independent-key attacks, where both use list decoding, is outside the scope of this paper. On the other hand, it is possible to compare independent-key attacks that use list decoding with related-key attacks that are based on uniquely decoded single-key attacks. As the value of ν of the former grows without bound if the list is small enough, and the value of ν of the latter can be kept constant (Theorem 12), Corollaries 13 and 14 also hold if the independent-key attack uses lists of estimates and the related-key attack is based on unique estimates; however improvements in the value of ν will be smaller for the same values of k , n and ε .

5 Experimental Verification

Corollary 14 says that any required fractional improvement over the cost of a single-key attack is possible if ψ is chosen well, if ε is small enough, and if n and k are large enough. In this section we examine how much improvement is possible with small values of k and n for linear cryptanalysis on reduced-round DES. We also examine the validity of the main unvalidated assumption of our theoretical results, that individual uses of the super-channel, for different keys, are independent of the relationship among the keys.

Motivated by Forney’s classical constructions of concatenated codes, we use Reed-Solomon (RS) codes as the relationship among the keys. We use Matsui’s linear cryptanalytic attack [16] as the statistical single-key recovery attack; the results of [16] have not been substantially improved upon, and the description of the attack is complete enough to allow for a correct reproduction. Our experimental results demonstrate that, using related-key recovery, we can obtain a reasonable improvement in ν over that of independent-key recovery, for linear cryptanalysis on reduced-round DES and small values of n and k . We also find that the results are close to those predicted by Reed-Solomon code error expressions, hence the assumption that individual uses of the super-channel are independent of the key is valid. (Because this assumption was unvalidated before the experiments were carried out, we were not able to simply predict the error of the related-key attacks using known error expressions for Reed-Solomon codes.)

5.1 Experimental Procedure

We first carry out linear cryptanalysis on 8-round DES, as described in [16], and verify that we are able to produce very similar results. We then calibrate the error of the single-key attack, ϵ , as a function of the number of P/C pairs, N . This provides the error of the super-channel. Finally, we carry out a number of related-key recovery attacks using the Reed-Solomon code for ψ , and various values of k and n . We compare the values of ν for the related and independent-key attacks.

5.2 Calibration of the Single-Key Recovery Attack

Matsui uses $(r - 1)$ -round linear cryptanalysis to determine twelve key bits, using a single linear approximation of $r - 1$ rounds, with a bias of 1.95×2^{-9} . We perform the complete experiment

for 8-round DES described in [15], with three minor differences. First, we use the key schedule of DES (where some key bits are related), whereas [15] uses a schedule where key bits are not related. We do not use the relationship among the key bits to improve our estimate, hence we do not expect the relationship to affect our results – this is corroborated by the fact that our results are very similar to those of [15]. Second, the results in [15] were performed for 8-round DES, and presented as predictions for full DES. Thus, for example, the well-known prediction of an attack requiring 2^{43} P/C pairs for 16-round DES corresponds to an experimental result that uses 1.49×2^{17} (approximately, 200,000) P/C pairs for 8-round DES. We too perform the results for 8-round DES, and present them as such. Finally, [15] used 100,000 instances of the attack to characterize the accuracy as a function of N , we use 10,000. Note that [15] describes an attack on full DES, as well as the possibility of determining a total of twenty-six bits of the key through linear cryptanalysis, and all other bits through exhaustive search. However, the detailed experimental results described are for the determination of twelve key bits for 8-round DES.

We are able to reproduce Matsui’s original attack with minor variations attributable to the minor differences in our experimental procedures; for reasons of space, the results verifying our reproduction of Matsui’s results are not made available here; however, all experimental results may be found in [19]. The probability of error, ϵ , for unique key-recovery, as a function of the number of P/C pairs, N , provides the probability of error of the super-channel for a given codeword length. If the adversary were able to transmit at channel capacity, N would be $\frac{12}{\mathcal{C}}$, (where \mathcal{C} is the capacity of the cryptanalytic channel, which may be approximated as $\frac{2(1.95 \times 2^{-9})^2}{\ln 2}$), about 287,000 P/C pairs, for any probability of error. In Matsui’s original experiments, about 200,000 P/C pairs resulted in an error of greater than 50% in unique-key recovery [15].

5.3 Limits on Related-Key Recovery

Recall from Theorem 12 that n related keys constructed from k independent keys can be used to obtain a related-key attack with amortized cost as low as $\frac{N}{d\mathcal{C}_S(N)}$. Figure 2 provides a plot of $\frac{N}{d\mathcal{C}_S(\epsilon)}$ in P/C pairs per key as a function of N . We observe that its minimum value (the tight bound of Corollary 15) is 535,000 P/C pairs per key, which is about 1.86 times the cost corresponding to the capacity of the basic cryptanalytic channel (about 287,000 P/C pairs per key, the bound of Corollary 16).

In the next section we describe the related-key recovery attacks we carried out.

5.4 Related-Key Recovery on Reduced-Round DES

The experiment performed is as follows. k independent keys, $(\ddot{l}_1, \ddot{l}_2, \dots, \ddot{l}_k)$, are chosen. These are encoded using an (n, k) Reed Solomon code to obtain n related keys, $(\ddot{k}_1, \ddot{k}_2, \dots, \ddot{k}_n)$. The bits to be estimated by the attack are encoded independently of the other bits. That is, $\kappa(\ddot{l}_1), \kappa(\ddot{l}_2), \dots, \kappa(\ddot{l}_k)$, are encoded to obtain $\kappa(\ddot{k}_1), \kappa(\ddot{k}_2), \dots, \kappa(\ddot{k}_n)$, and the values of the other forty-four bits of each key are encoded separately. n single-key-recovery attacks are carried out using $\ddot{k}_1, \ddot{k}_2, \dots, \ddot{k}_n$. The estimates obtained, $\widehat{\kappa(\ddot{k}_1)}, \widehat{\kappa(\ddot{k}_2)}, \dots, \widehat{\kappa(\ddot{k}_n)}$ are RS-decoded to obtain $\widehat{\kappa(\ddot{l}_1)}, \widehat{\kappa(\ddot{l}_2)}, \dots, \widehat{\kappa(\ddot{l}_k)}$. 1000 uniformly random instances of the experiment are carried out for each value of (n, k) . The rate of the outer code was maintained at approximately $\frac{2}{3}$, and the values of (n, k) used were: (7, 5), (15, 11), (31, 21), (63, 43), and (127, 87). The experimental key recovery errors were compared to the errors predicted by the Reed-Solomon decoding error formulae for given channel error probability [6], and found to be very close.

In Figure 3 we provide several plots. One plot corresponds to the single key recovery attack.

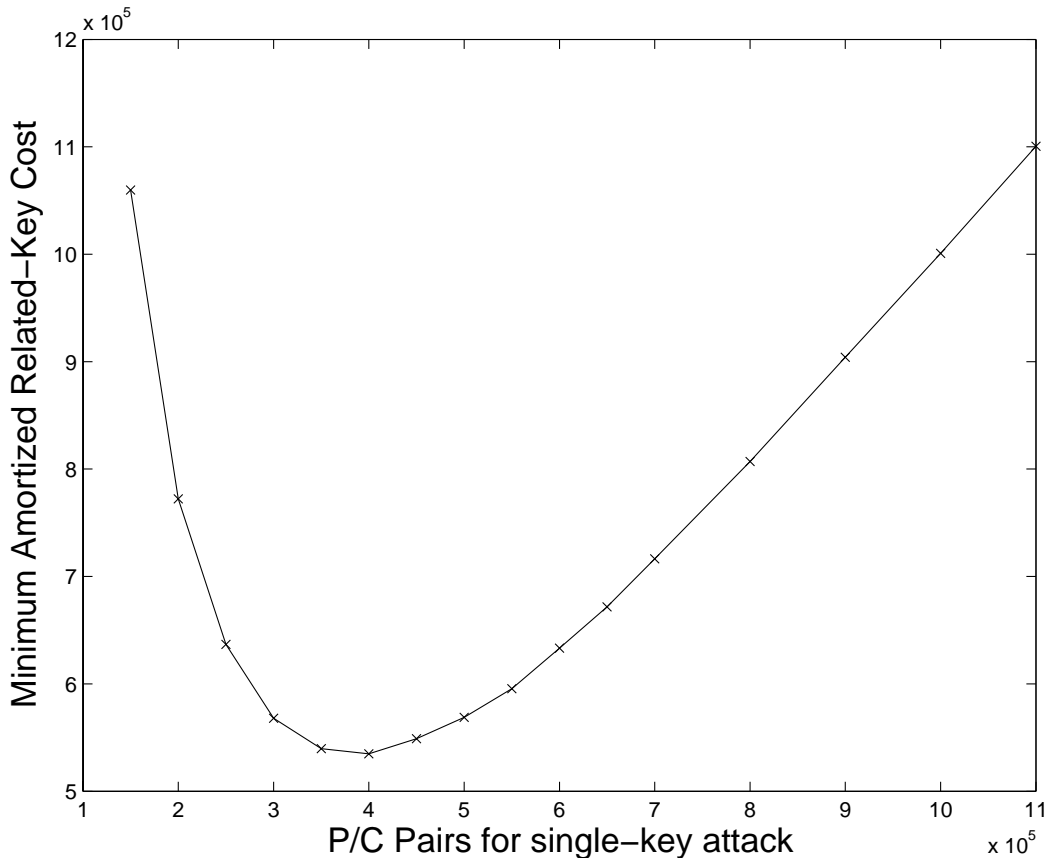


Figure 2: Minimum value of ν , $\frac{N}{dC_S(\epsilon)}$

Each of the other plots corresponds to a particular value of (n, k) . The x-axis provides the value of ν in P/C pairs per independent key. The y-axis provides the value of ϵ for an independent-key attack using k independent keys that would result in the same value of ϵ as the (n, k) related-key attack. That is, the y-axis provides the value of $1 - (1 - \epsilon)^{\frac{1}{k}}$. Thus, the solid line curve represents the independent-key attack for any value of k . If the y-axis value for an (n, k) attack is lower than that of the solid line, the related-key attack is more efficient than the independent-key attack. For example, the value of ϵ required to obtain the performance of the related-key attack with $(n, k) = (127, 87)$ is about 0.02, while that of the single-key attack is 0.04, for $\nu = 5.5 \times 10^5$ P/C pairs per independent key. Thus the $(127, 87)$ RS code provides a better attack at this error probability. We observe that the $(7, 5)$ RS code offers improvements at very low error levels, while the $(127, 87)$ code offers improvements at higher error levels, as expected.

Note that the values of ϵ beyond which it is more efficient to use the corresponding RS-encoded related-key attack is monotonic decreasing as a function of n , as expected.

Finally, we observed that the related-key attack has an 11% lower value of ν than the independent-key attack, if $\epsilon = 0.04$ and $k = 11$, a 22% lower value if $\epsilon = 0.03$ and $k = 21$, and a 27% lower value if $\epsilon = 0.17$ and $k = 87$. The difference between the independent and related-key attacks is further illustrated by the fact that the values of ϵ with the same smaller value of ν when $k = 87$ are 0.17 for the related-key attack, and 0.81 for the independent-key attack.

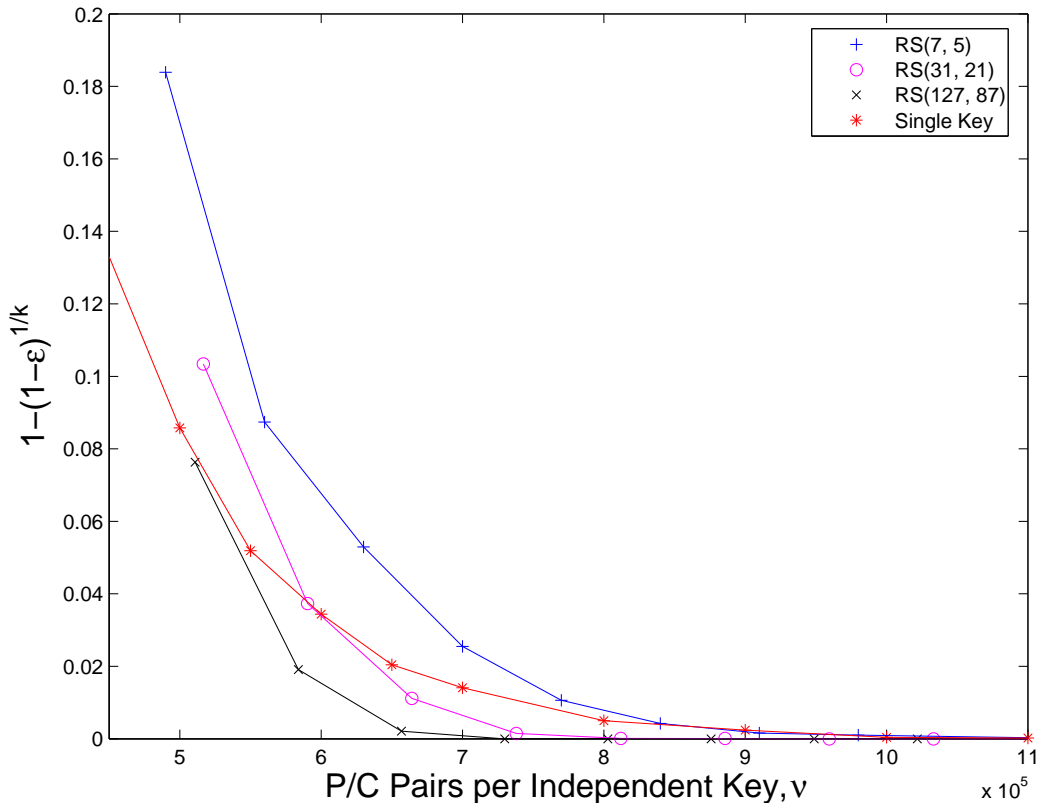


Figure 3: Equivalent single-key recovery error as a function of ν

6 Conclusions and Future Directions

We have presented a definition of the known-plaintext statistical key-recovery attack on a block cipher and demonstrated that it is a generalization of several common attacks. We have also presented a Cryptanalytic Channel Model (CCM) that treats the cipher as a channel, and statistical key recovery as communication of an encoded value of the key over the channel. We have examined how a relationship among keys affects the error of key recovery, and have found that related-key recovery attacks can asymptotically achieve lower amortized cost than an equivalent set of many single-key attacks. This result does not depend on specific properties of the cipher, but simply on the fact that it is vulnerable to statistical cryptanalysis. Finally, we have presented experimental results to support the model and to demonstrate the extent to which our asymptotic results are applicable for a small number of independent and related keys.

A number of future directions present themselves. First an examination of related-key recovery within the list decoding framework [10] might result in more efficient attacks, and could also provide insights into what types of round functions are resilient to such attacks. Second, the ideas of this paper can also be applied to keys that are not related in a deterministic fashion, but when there is a weaker (probabilistic) relationship among the keys. Third, an examination of key scheduling algorithms in this framework could be very interesting. Finally, other attacks, such as ciphertext-only attacks, are also expected to lend themselves well to study in this framework.

7 Acknowledgements

The authors would like to acknowledge suggestions by, and discussions with, David Wagner.

References

- [1] Mihir Bellare and Tadayoshi Kohno. A theoretical treatment of related-key attacks: Rka-prps, rka-prfs, and applications. In *EUROCRYPT '03: Workshop on the theory and application of cryptographic techniques on Advances in cryptology*, 2003.
- [2] Eli Biham. New types of cryptanalytic attacks using related keys. In *EUROCRYPT '93: Workshop on the theory and application of cryptographic techniques on Advances in cryptology*, 1993.
- [3] Eli Biham and Adi Shamir. Differential cryptanalysis of des-like cryptosystems. In *CRYPTO '90: Proceedings of the 10th Annual International Cryptology Conference on Advances in Cryptology*, 1991.
- [4] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*, chapter 8. Wiley-Interscience, 1991.
- [5] Eric Filiol. Plaintext-dependant repetition codes cryptanalysis of block ciphers - the aes case. Cryptology ePrint Archive, Report 2003/003, 2003. <http://eprint.iacr.org/>.
- [6] David G. Forney. *Concatenated Codes*. MIT Press, 1966.
- [7] Carlo Harpes, Gerhard G. Kramer, and James L. Massey. A generalization of linear cryptanalysis and the applicability of matsui's piling-up lemma. In *EUROCRYPT '95: Proceedings of the International Conference on the Theory and Applications of Cryptographic Techniques*, 1995.
- [8] Thomas Jakobsen. Correlation attacks on block ciphers. Master's thesis, Dept. of Mathematics, Technical University of Denmark,, 1996.
- [9] Thomas Jakobsen. Cryptanalysis of block ciphers with probabilistic non-linear relations of low degree. In *CRYPTO '98: Proceedings of the 18th Annual International Cryptology Conference on Advances in Cryptology*, 1998.
- [10] Thomas Jakobsen. *Higher-Order Cryptanalysis of Block Ciphers*. PhD thesis, Dept. of Mathematics, Technical University of Denmark, 1999.
- [11] Thomas Jakobsen and Lars R. Knudsen. The interpolation attack on block ciphers. In *FSE '97: Proceedings of the 4th International Workshop on Fast Software Encryption*, 1997.
- [12] Thomas Jakobsen and Lars R. Knudsen. Attacks on block ciphers of low algebraic degree. *Journal of Cryptology*, 14(3):197–210, 2001.
- [13] John Kelsey, Bruce Schneier, and David Wagner. Key-schedule cryptanalysis of idea, g-des, gost, safer, and triple-des. In *CRYPTO '96: Proceedings of the 16th Annual International Cryptology Conference on Advances in Cryptology*, 1996.
- [14] John Kelsey, Bruce Schneier, and David Wagner. Related-key cryptanalysis of 3-way, biham-des, cast, des-x, newdes, rc2, and tea. In *ICICS '97: Proceedings of the First International Conference on Information and Communication Security*, 1997.
- [15] Mitsuru Matsui. The first experimental cryptanalysis of the data encryption standard. In *CRYPTO '94: Proceedings of the 18th Annual International Cryptology Conference on Advances in Cryptology*, 1994.
- [16] Mitsuru Matsui. Linear cryptanalysis method for des cipher. In *EUROCRYPT '93: Workshop on the theory and application of cryptographic techniques on Advances in cryptology*, 1994.
- [17] W. Meier and O. Staffelbach. Fast correlation attacks on stream ciphers. In *EUROCRYPT '88: Workshop on the theory and application of cryptographic techniques on Advances in cryptology*, 1988.

- [18] W. Meier and O. Staffelbach. Fast correlation attack on certain stream ciphers. *Journal of Cryptology*, pages 159–176, 1989.
- [19] Darakhshan J. Mir. Related-key linear cryptanalysis of des. Master’s thesis, School of Engineering and Applied Science, George Washington University, 2006.
- [20] S. Murphy, F. Piper, M. Walker, and P. Wild. Maximum likelihood estimation for block cipher keys, 1994.
- [21] Claude Shannon. A mathematical theory of communication. *Bell Systems Technical Journal*, 27:379–423, 1948.
- [22] Harry L. Van Trees. *Detection, Estimation, and Modulation Theory, Part I*. John Wiley & Sons, 1968.
- [23] S. Vaudenay. Decorrelation: A theory for block cipher security. *Journal of Cryptology*, 16(4):249–286, Sept. 2003.
- [24] Serge Vaudenay. An experiment on des statistical cryptanalysis. In *CCS ’96: Proceedings of the 3rd ACM conference on Computer and Communications Security*, 1996.
- [25] Poorvi L. Vora. An information-theoretic approach to inference attacks on random data perturbation and a related privacy measure. *IEEE Trans. Info. Theory*, 53(8), 2007. 2971–2977.
- [26] Poorvi L. Vora and Darakhshan J. Mir. Related-key linear cryptanalysis. In *ISIT ’06: Proceedings of the 2006 IEEE International Symposium of Information Theory*, pages 1609–1613, July 2006.
- [27] David Wagner. Towards a unifying view of block cipher cryptanalysis. In *FSE ’04: Eleventh International Workshop on Fast Software Encryption*, 2004.