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,
George Washington University

June 11, 2007

Abstract

This paper presents the *Cryptanalytic Channel Model (CCM)*. The model treats statistical key recovery as communication over a low capacity channel, where the channel and the encoding are determined by the cipher and the specific attack. A new attack, related-key recovery – the use of n related keys generated from k independent ones – is defined for all ciphers vulnerable to single-key recovery. It is shown to correspond to the use of a concatenated code over the channel, where the relationship among the keys determines the outer code, and the cipher and the attack the inner code. It is shown that there exists a relationship among keys for which the communication complexity per bit of independent key is finite, *for any probability of key recovery error*. This may be compared to the unbounded communication complexity per bit of the single-key-recovery attack. The practical implications of this result are demonstrated through experiments on reduced-round DES.

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

Communicating Author: Poorvi L. Vora, 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 [11, 21, 25] 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.

The problem may be described more formally as follows. Consider a known-plaintext statistical key-recovery attack on a block cipher. The attack is defined formally in section 3.1, and is referred to as *single key recovery*; examples include linear [16], differential [3], integral [14] or noisy polynomial [10] cryptanalysis. Suppose the attack uses N plaintext-ciphertext (P/C) pairs to determine d bits of a single key. As with other theoretical analyses of statistical cryptanalysis [16, 18], it is assumed that the attacker obtains a unique estimate of the d bits. (Section 4.4 describes how the ideas contained in this paper may be applied to the case where the attacker obtains a small list of estimates of the key, instead of a unique estimate). Denote by ϵ the corresponding probability of error in key recovery. It is well-known that as ϵ approaches zero, N increases without bound – see, for example, [16, Lemmas 2 and 5]. Because d is fixed, the communication complexity per bit, $\frac{N}{d}$, also increases indefinitely. That is, if ν is the communication complexity per bit of key determined,

$$\epsilon \to 0 \Rightarrow N \to \infty \Rightarrow \nu = \frac{N}{d} \to \infty \tag{1}$$

Now consider the case when multiple keys are used, and the adversary is able to obtain N P/C pairs for each of k independent keys. This attack is defined formally in section 4, and is referred to as *independent-key recovery*. If ε (defined more precisely in section 4) represents the probability that any of the k estimated keys is in error, a relationship such as that of (1) still holds (also see Lemma 10):

$$\varepsilon = 1 - (1-\epsilon)^k \to 0 \Rightarrow \epsilon \to 0 \Rightarrow N \to \infty \Rightarrow \nu = \frac{kN}{kd} \to \infty \tag{2}$$

Finally, consider the related-key attack, also defined in section 4, and referred to as *related-key recovery*. Suppose n ($n > k$) related keys are constructed from k independent ones. Suppose the adversary obtains M P/C pairs for each related key, $M < N$ (that is, the related keys are changed "more frequently" than the independent ones). The d bits of each of the n related keys are determined in independent single-key-recovery attacks, after which the relationship among the keys is used to determine the corresponding bits of the k independent keys. How does the value of ν for related-key recovery, $\frac{nM}{kd}$, compare with that of independent-key recovery, $\frac{N}{d}$, for the same value of ε?

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. Linear, differential, integral and noisy polynomial cryptanalysis are shown to be examples of the attack and of channel communication.

- 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 that there exists a relationship among the keys such that:

$$\varepsilon \to 0 \text{ and } \nu \simeq \Lambda \tag{3}$$

    for some constant finite Λ. (Compare (3) to (2)). Further, Λ is also shown to be asymptotically bounded below.
  - As a corollary, of (2) and (3), it shows 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.
- It provides experimental results for related-key linear cryptanalysis on reduced-round DES. For example, 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.

Thus this paper demonstrates that 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. This is similar to the advantage of using channel codes over repetition codes for channel communication. It also demonstrates a limit on the adversary: if $\varepsilon \to 0$, ν is asymptotically bounded below. This is similar to the upper bound of channel capacity on the rate of a channel code. 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 $N \to \infty$ as $\epsilon \to 0$. 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 [24], 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 [17].

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)*. It also shows how the model applies to several examples of common single-key-recovery attacks. 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

Filiol [5] first suggested that a known probabilistic relationship – between ciphertext and a single binary property of the key – be modeled as a communication channel. In his model, for ciphertext-only attacks, 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). [5] 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.

In other related work, Jakobsen [11] treats attacks on ciphers whose properties can be modeled as polynomials of small degree, and uses recent work in computational coding theory to efficiently decode attacks. In particular, he proposes the list decoding model, where the key estimate consists of a small set of possibilities, as opposed to a unique estimate.

The framework of Wagner [25] describes the techniques for obtaining the probabilistic relationships among the plaintext, ciphertext and key. It models the relationships as Markov chains, in the manner of [21, 22].

Biham examines related-key attacks on block ciphers, tracing the relationships among the keys to the key scheduling algorithm [2]. Kelsey, Schneier and Wagner [12, 13] present related-key attacks on various block ciphers, and demonstrate how real protocols can be exploited to mount such attacks.

No single framework prior to the CCM addresses both related-key recovery and statistical cryptanalysis. The CCM extends the model of [5] to include known-plaintext attacks and related-key attacks. While [5] 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 [21, 25], the CCM models the relationship among plaintext and ciphertext as a communication channel, and not as a Markov chain. This allows the CCM to address related-key attacks, and also allows access to a rich literature in coding theory. At the same time, the CCM allows, in a very natural way, the use of [21, 25] 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 $\ddot{K}$, 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 $\ddot{K} = \ddot{k}$. These are generated by

picking x_j uniformly at random and encrypting it using the block cipher and key $\ddot{k}$ 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 set of plaintext values, $\mathbf{X}$, and an associated set 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 $\ddot{k}$. 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 bits, assuming, as is typical in estimation theory (see, for example, [20]), a uniform *a priori* distribution on $\ddot{K}$. Before formally defining the attack, we present an example.

Example 1: Linear Cryptanalysis using an r-round Approximation

Consider r-round linear cryptanalysis on an r-round iterated cipher [16, 8]. The cipher is approximated using a linear (or affine) expression:

$$Pr[f_r(X) \oplus g_r(Y) = h_r(\ddot{k})] = \frac{1}{2} + \gamma \tag{4}$$

where f_r, g_r and h_r denote linear or affine functions, and γ the bias; all are independent of fixed key $\ddot{k}$. The probability is taken over all possible plaintexts. In this case, $S = f_r(X) \oplus g_r(Y)$, and $T = h_r(\ddot{k})$. The adversary uses N P/C pairs: $\{(x_j, y_j)\}_{j=1}^N$ to determine the maximum-likelihood estimate of $h_r(\ddot{k})$, denoted $\widehat{h_r(\ddot{k})}$. It is the one that satisfies (4) most often. That is,

$$\widehat{h_r(\ddot{k})} = \overset{argmax}{z} |\{(x_j, y_j) | f_r(x_j) \oplus g_r(y_j) = z\}|$$

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 $\ddot{k} \in \mathcal{K}$ consists of:*

- *Function κ, $\kappa : \mathcal{K} \to \kappa(\mathcal{K}) \subseteq \mathcal{K}$*
- *A function* Many *mapping a single plaintext X to a set of plaintexts $\mathbf{X}$,* Many$(X) = \mathbf{X}$*. The corresponding set of ciphertexts is denoted $\mathbf{Y}$*
- *Random variables $S(\mathbf{X}, \mathbf{Y}) \in \mathcal{Z}$ and $T(\mathbf{X}, \mathbf{Y}, \kappa(\ddot{k})) \in \mathcal{Z}$, where $\mathcal{Z}$ denotes some domain, its size is denoted 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 $|.|$ 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 $Z \in \mathcal{Z}$:*

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

for small positive γ.

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

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

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, while S may play the part of a statistical distinguisher[1] – for example when T is independent of $\mathbf{X}$ and $\mathbf{Y}$ – 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 $\mathcal{K}$.

The attack of example 1 satisfies Definition 1, with

$$\begin{aligned} \mathcal{Z} &= \mathbb{Z}_2 \\ \mathbf{X} &= X \\ \kappa &= h_r \\ d &= 1 \\ S(X,Y) &= f_r(X) \oplus g_r(Y) \\ T(X,Y,\kappa(\ddot{k})) &= \kappa(\ddot{k}) \end{aligned}$$

We now define the key-recovery error.

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

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

For the attack of example 1, it is known that [16, Lemmas 2 and 5]:

$$\epsilon \to 0 \Rightarrow N \to \infty \tag{5}$$

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). Before we describe the CCM, we briefly review the notion of channel communication, and observe how the attack of example 1 corresponds to channel communication.

[1]For our purposes, a statistical distinguisher is 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.

A communication channel is a triplet, consisting of an input alphabet, an output alphabet, and the conditional probability distribution of the output given the input [4]. For the typical channel, the input and output alphabets are identical, and the probability of error is the probability that the output is not equal to the input. For a simplex channel, each incorrect output is equally likely. The capacity of the channel, $\mathcal{C}$, is the maximum value, over all input probability distributions, of the mutual information between channel output and input.

A message transmitted across a channel typically includes redundancy, which enables error correction. Thus, before transmission, a message is encoded using a channel code. An (n, k) channel code is an injective function that maps a message of length k to a codeword of length n, such that $n \geq k$. (A simple example of the channel code is the $(n, 1)$ repetition code, where a codeword consists of the message symbol repeated n times.) The codeword is transmitted, and, after the probabilistic perturbation of the channel, received at the output. The received codeword is decoded to obtain either a single estimate of the message (known as unique decoding) or a small list of estimates (known as list decoding). The communication efficiency of the codeword is measured by its rate $R = \frac{k}{n}$. For example, the rate of the repetition code is $\frac{1}{n}$. In general, a lower decoding error is possible with a lower rate. In the case of the repetition code, the error decreases exponentially with n, and a decrease in error may only be brought about with a corresponding decrease in rate.

Example 2: Linear Cryptanalysis using an r-round Approximation as Channel Communication

The attack of example 1 may be modeled as channel communication as follows (see Figure 1). The input to the channel is the value of $h_r(\ddot{k})$, which is not directly accessible to the adversary. The adversary can, however, compute $f_r(X) \oplus g_r(Y)$ from plaintext and ciphertext, which provides the value of $h_r(\ddot{k})$ slightly more often than not – see (4). Hence $f_r(X) \oplus g_r(Y)$ is the channel output. The input and output alphabets are identical, $\mathbb{Z}_2$. The randomness of the channel is provided by the different values of plaintext encrypted, and the probability of channel error is $p_e = \frac{1}{2} - \gamma$; this is assumed to be independent of the value of the key and hence of $h_r(\ddot{k})$ (that is, the channel is a simplex channel). $\mathcal{C} \simeq \frac{2\gamma^2}{ln2}$.

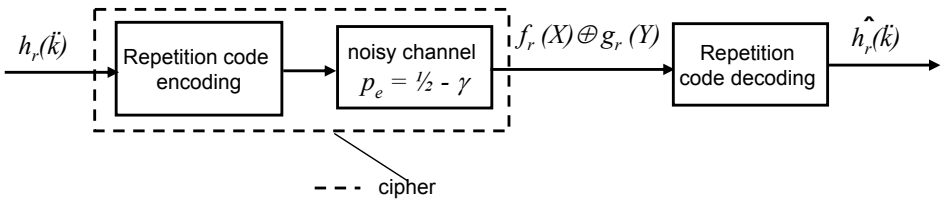

Figure 1: CCM: Linear Cryptanalysis – r Round Approximations

For N P/C pairs, the cipher transmits the value of $h_r(\ddot{k})$ over the channel N times – that is, it transmits a length-N repetition code over $h_r(\ddot{k})$. The procedure of estimating $h_r(\ddot{k})$ from $\{f_r(x_j) \oplus g_r(y_j)\}_{j=1}^{N}$ is the (maximum-likelihood) decoding procedure for a length-N repetition code. The rate of transmission is $\frac{1}{N}$, or the inverse of the communication complexity per bit of key determined. (5) is the equivalent of the fact that a decrease in repetition code decoding error may only be obtained through a decrease in rate.

We now turn to the problem of representing the statistical attack as channel communication.

First, we define the *cryptanalytic channel* over which the communication occurs.

Definition 3 *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* $\mathcal{Z}$*.*

$\kappa(\ddot{k})$ forms the message (see Figure 2). 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 $[T(\mathbf{X}_1, \mathbf{Y}_1, \kappa(\ddot{k})), T(\mathbf{X}_2, \mathbf{Y}_2, \kappa(\ddot{k})), ...T(\mathbf{X}_N, \mathbf{Y}_N, \kappa(\ddot{k}))]$ forms the transmitted codeword, of length N. $[S(\mathbf{X}_1, \mathbf{Y}_1)$, $S(\mathbf{X}_2, \mathbf{Y}_2)$, ... $S(\mathbf{X}_N, \mathbf{Y}_N)]$ forms the received codeword. The decoding algorithm is Algorithm $KeyRecovery$, and corresponds to maximum-likelihood decoding. R is the inverse of the communication complexity per key bit, which is $\frac{N}{d}$. The capacity of the channel, for small γ, is approximately $\mathcal{C} = \frac{m^2\gamma^2}{2(m-1)ln2}$.

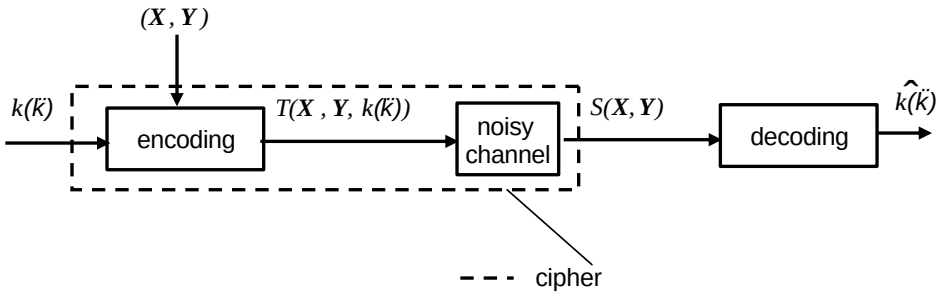

Figure 2: The Cryptanalytic Channel Model

Definition 1 does not 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. In section 3.4, we show that differential and integral key recovery attacks are typically of this kind, and the more efficient polynomial and linear cryptanalytic attacks are also of this kind.

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

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 [9]. 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 $\ddot{k}_{-r}$. Hence:

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

for $\gamma > 0$, and some function κ_{-r}. Observing that $Y^{(r-1)} = \mathcal{F}^{-1}_{\ddot{k}^{(r)}}(Y)$, 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:

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

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

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

One has the following definition.

Definition 4 *A known-plaintext statistical single-key recovery attack* Γ *on an iterated block cipher is a* $k^{(r)}$-recovery *attack if* $\exists$

- $\kappa_{-r} : \mathcal{K}_{-r} \rightarrow \kappa_{-r}(\mathcal{K}_{-r}) \subseteq \mathcal{K}_{-r}$*, where* $\mathcal{K}_{-r}$ *is the space of all keys for the* $(r-1)$*-round cipher.*
- *r.v.* $S'(\mathbf{X}) \in \mathcal{Z}$
- *r.v.* $T'(\mathbf{X}, \mathcal{F}^{-1}_{\ddot{k}^{(r)}}(\mathbf{Y}), \kappa_{-r}(\ddot{k}_{-r})) \in \mathcal{Z}$

such that:

- $S(\mathbf{X}, \mathbf{Y}) = S'(\mathbf{X})$
- $T(\mathbf{X}, \mathbf{Y}, \kappa(\ddot{k})) = T'(\mathbf{X}, \mathcal{F}^{-1}_{\ddot{k}^{(r)}}(\mathbf{Y}), \kappa_{-r}(\ddot{k}_{-r}))$

Note that, if T' is constant as a function of $\ddot{k}_{-r}$ (such as in integral and differential cryptanalysis, see examples 3.4 and 3.4), κ_{-r} is a trivial function. Note also that S' may be constant as a function of $\mathbf{X}$.

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 $\mathcal{I}_j$:

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

The codeword, α, is:

$$\alpha = [\mathcal{I}_1(\kappa_{-r}(\ddot{k}_{-r}), \ddot{k}^{(r)}), \mathcal{I}_2(\kappa_{-r}(\ddot{k}_{-r}), \ddot{k}^{(r)}), ...\mathcal{I}_N(\kappa_{-r}(\ddot{k}_{-r}), \ddot{k}^{(r)})]$$

The received codeword is

$$[S'(\mathbf{x}_1), S'(\mathbf{x}_2), ...S'(\mathbf{x}_N)]$$

3.4 Some Examples

In this section, we provide a few examples of $\ddot{k}^{(r)}$ recovery.

Example 3: Linear Cryptanalysis: $(r-1)$-round approximation

In practice, the attack of example 1 is not very efficient, and several P/C pairs provide only one bit of the key. An improvement on the attack [15] is a $k^{(r)}$-recovery attack that uses 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:

$$Pr[f_{r-1}(X) = g_{r-1}(\mathcal{F}^{-1}_{\ddot{k}^{(r)}}(Y)) \oplus h_{r-1}(\ddot{k}_{-r})] = \frac{1}{2} + \gamma \tag{6}$$

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 $\phi(z) = |\{(x_j, y_j) \ : \ f_{r-1}(x_j) = g_{r-1}(\mathcal{F}_z^{-1}(y_j))\}|$,

$$\widehat{\ddot{k}^{(r)}} = \stackrel{argmax}{z} \quad || \ \phi(z) \ - \ \frac{N}{2} \ ||$$

where $||.||$ denotes the absolute value. If $\phi(\widehat{\ddot{k}^{(r)}}) > \frac{N}{2}$, $\widehat{h_{r-1}(\ddot{k}_{-r})} = 0$, else $\widehat{h_{r-1}(\ddot{k}_{-r})} = 1$. 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 Definitions 1 and 4, with

$$\begin{aligned}
\mathcal{Z} &= \mathbb{Z}_2 \\
\mathbf{X} &= X \\
\kappa_{-r} &= h_{r-1} \\
d &= \frac{b}{r} + 1 \\
A(X, Y^{(r-1)}) &= f_{r-1}(X) \oplus g_{r-1}(Y^{(r-1)}) \\
B(\mathbf{X}, \mathbf{Y}^{(r-1)}, \kappa_{-r}(\ddot{k}_{-r})) &= h_{r-1}(\ddot{k}_{-r}) \\
S'(X) &= f_{r-1}(X) \\
T'(X, \mathcal{F}_{\ddot{k}^{(r)}}^{-1}(Y), \kappa_{-r}(\ddot{k}_{-r})) &= g_{r-1}(\mathcal{F}_{\ddot{k}^{(r)}}^{-1}(Y)) \oplus h_{r-1}(\ddot{k}_{-r})
\end{aligned}$$

Example 4: Linear Cryptanalysis: $(r-1)$-round approximation as Channel Communication

The attack of example 3 is different from that of examples 1 and 2 in two ways: the bias is strictly greater than the bias of the previous attack, and the codeword transmitted is *not the repetition code* on a single bit of the key, but a ciphertext-dependent code on $\frac{b}{r}+1$ key bits. 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 (6)). 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 $\ddot{k}_{-r}$.

As in example 2, the codeword itself is not accessible to the adversary. However, $[(f_{r-1}(x_1), f_{r-1}(x_2), ...f_{r-1}(x_N))]$ is a very noisy value of the codeword, providing the output of the cryptanalytic channel (see Figure 3 and (6)). The channel error probability is $p_e = \frac{1}{2} - \gamma$, the channel is symmetric, and the corresponding capacity is $\mathcal{C} \simeq \frac{2\gamma^2}{ln2}$.

$\ddot{k}^{(r)}$ and $h_{r-1}(\ddot{k}_{-r})$ are determined from the values of $f_{r-1}(x_j)$ using maximum-likelihood decoding [18]. While the computational complexity of this attack is greater than that of the one of example 1, as it performs an exhaustive search for $\ddot{k}^{(r)}$, experiments typically indicate that the communication complexity per bit of key determined, $\frac{N}{\frac{b}{r}+1}$, is considerably smaller [16] – this is because the channel has larger capacity, and the encoding is a better code.

Example 5: Noisy Polynomial Interpolation Attacks

Interpolation attacks [10] approximate $Y^{(r-1)}$, the output of $r-1$ rounds of the cipher, as a low-degree polynomial in X, treating both X and $Y^{(r-1)}$ as belonging to a finite field (such as,

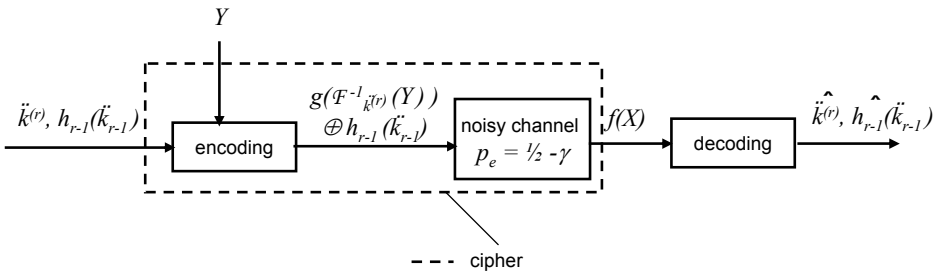

Figure 3: CCM: Linear Cryptanalysis – $r-1$ Round Approximations

for example, $GF(2^q)$). If the polynomial is $p(x) = \sum_{i=1}^{m} a_i x^i$, the coefficients $\mathbf{a}_r = (a_1, a_2, ...a_m)$ depend on $\ddot{k}_{-r}$, and

$$Pr[0 = \mathcal{F}^{-1}_{\ddot{k}^{(r)}}(Y) - \sum_{i=1}^{m} a_i X^i] = \frac{1}{2^n} + \gamma \tag{7}$$

for some bias γ. Noisy polynomial interpolation cryptanalysis satisfies Definitions 1 and 4, with

$$\begin{aligned}
\mathcal{Z} &= GF(2^q) \\
\mathbf{X} &= X \\
\kappa_{-r}(\ddot{k}_{-r}) &= \mathbf{a}_{r-1} \\
A(X, Y^{(r-1)}) &= Y^{(r-1)} \\
B(X, Y^{(r-1)}, \kappa_{-r}(\ddot{k}_{-r})) &= \sum_{i=1}^{m} a_i X^i \\
S'(X) &= 0 \\
T'(X, \mathcal{F}^{-1}_{\ddot{k}^{(r)}}(Y), \kappa_{-r}(\ddot{k}_{-r})) &= \sum a_i X^i - \mathcal{F}^{-1}_{\ddot{k}^{(r)}}(Y)
\end{aligned}$$

Example 6: Noisy Polynomial Cryptanalysis as Channel Communication

In the CCM for noisy polynomial interpolation, the message consists of $\ddot{k}^{(r)}$ and $\mathbf{a}$. The codeword consists of symbols in $GF(2^q)$, the j^{th} code symbol is $\mathcal{I}_j(\ddot{k}^{(r)}, \mathbf{a}) = \sum a_{r,i} x_j^i - \mathcal{F}^{-1}_{\ddot{k}^{(r)}}(y_j)$, see (7). Notice that the codeword is plaintext and ciphertext dependent. The received symbol is always the zero symbol, and this is a noisy version of the transmitted codeword.

Example 7: 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 $X_1 \oplus X_2 = \Delta X$,

$$Pr[\Delta Y = \mathcal{F}^{-1}_{\ddot{k}^{(r)}}(Y_1) \oplus \mathcal{F}^{-1}_{\ddot{k}^{(r)}}(Y_2)] = \frac{1}{2^q} + \gamma \tag{8}$$

for some $\gamma > 0$, as $Y_i^{(r-1)} = \mathcal{F}^{-1}_{\ddot{k}^{(r)}}(Y_i)$, $i = 1, 2$. Differential cryptanalysis satisfies Definitions 1 and

4, with

$$\begin{aligned}
\mathcal{Z} &= \Sigma^q \\
\mathbf{X} &= \{X, X \oplus \Delta P\} \\
\kappa(\ddot{k}) &= \ddot{k}^{(r)} \\
d &= \frac{b}{r} \\
A(\mathbf{X}) &= \Delta Y \ (a\ constant) \\
B(\mathbf{X}, \mathbf{Y}^{(r-1)}, \kappa_{-r}(\ddot{k}_{-r})) &= Y_1^{(r-1)} \oplus Y_2^{(r-1)} \\
S'(\mathbf{X}) &= \Delta Y \\
T'(\mathbf{X}, \mathcal{F}_{\ddot{k}^{(r)}}^{-1}(Y), \kappa(\ddot{k})) &= \mathcal{F}_{\ddot{k}^{(r)}}^{-1}(Y_1) \oplus \mathcal{F}_{\ddot{k}^{(r)}}^{-1}(Y_2)
\end{aligned}$$

Example 8: Integral cryptanalysis

The integral cryptanalytic attack approximates $r-1$ rounds of the cipher so that a particular sum of plaintext values c results in a particular sum c' of corresponding ciphertext output with probability greater than $\frac{1}{2^q}$. If $\mathcal{D} \subset \Sigma^q$ is a set of plaintext such that $\sum_{v \in \mathcal{D}} v = c$ for some constant c and $E_{\ddot{k}}(\mathcal{D})$ the ciphertexts corresponding to plaintexts in $\mathcal{D}$:

$$Pr[c' = \sum_{w \in E_{\ddot{k}}(\mathcal{D})} \mathcal{F}_{\ddot{k}^{(r)}}^{-1}(w)] = \frac{1}{2^q} + \gamma \tag{9}$$

Integral cryptanalysis satisfies Definitions 1 and 4, with

$$\begin{aligned}
\mathcal{Z} &= \Sigma^q \\
\mathbf{X} &= \mathcal{D} \subset \Sigma^q \\
\kappa(\ddot{k}) &= \ddot{k}^{(r)} \\
d &= \frac{b}{r} \\
A(\mathbf{X}) &= c' \\
B(\mathbf{X}, \mathbf{Y}^{(r-1)}, \kappa_{-r}(\ddot{k}_{-r})) &= \sum_{y^{(r-1)} \in \mathbf{Y}^{(r-1)}} y^{(r-1)} \\
S'(\mathbf{X}) &= c' \\
T'(\mathbf{X}, \mathcal{F}_{\ddot{k}^{(r)}}^{-1}(\mathbf{Y}), \kappa(\ddot{k})) &= \sum_{y \in \mathbf{Y}} \mathcal{F}_{\ddot{k}^{(r)}}^{-1}(y)
\end{aligned}$$

Example 9: Differential and Integral Cryptanalysis as Channel Communication

In the CCM for differential and integral attacks, the message is $\ddot{k}^{(r)}$. The j^{th} symbol of the codeword is $\sum_{y \in \mathbf{y}_j} \mathcal{F}_{\ddot{k}^{(r)}}^{-1}(y)$. The received symbol is always ΔY for differential cryptanalysis, and c' for integral cryptanalysis, see (8) and (9). The decoding algorithm is maximum-likelihood decoding implemented through exhaustive search.

4 Related-Key Recovery

The repetition code, which consists of the transmission of a single symbol over the channel n times, requires a decrease in rate for a decrease in error. The transmission of related symbols, however, can provide a very different relationship between rate and error. Shannon's channel coding theorem

[19, 4] proves that there exists a channel code, using which the probability of decoding error can be made as small as desired, for constant 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, and 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,* $\ddot{\mathbf{l}} = (\ddot{l}_1, \ddot{l}_2, ...\ddot{l}_k)$*, such that* $\ddot{l}_i \in \mathcal{K}\ \forall\ i$
- n *related keys,* $\ddot{\mathbf{k}} = (\ddot{k}_1, \ddot{k}_2, ...\ddot{k}_n)$*, 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:*

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

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

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

A statistical related-key recovery attack is denoted Ψ = *(*Γ*,* k*,* n*,* ψ*, 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* $\kappa(\ddot{\mathbf{l}}) \neq \widehat{\kappa(\ddot{\mathbf{l}})}$.

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

A simple lemma describes the relationship between the value of $\varepsilon(N, k, k, I)$ and $\epsilon(N)$.

Lemma 10 *For a given single-key-recovery attack* Γ *with key-recovery error* $\epsilon(N)$ *and the corresponding independent-key-recovery attack,* (Γ, k, k, I, RelatedKeyRecovery)*, with related-key-recovery error* $\varepsilon(N, k, k, I)$

$$1 - \varepsilon(N, k, k, I) = (1 - \epsilon(N))^k$$

Proof. The left hand side is the probability that all the bits of all k independent keys are correctly estimated, so is the right hand side. ■

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 11 *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 $m_1m_2m_3...m_{n_{conc}} \in \mathcal{X}^{n_{conc}}$ 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 (see also Figure 4).

Lemma 12 *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 $\kappa(\ddot{\mathbf{l}})$ and a single codeword from the outer code is of the form $\kappa(\ddot{\mathbf{k}})$, 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, $\widehat{\kappa(\ddot{\mathbf{k}})}$. The procedure of finding $\widehat{\kappa(\ddot{\mathbf{l}})}$ from $\widehat{\kappa(\ddot{\mathbf{k}})}$ is the maximum likelihood decoding of ψ. ■

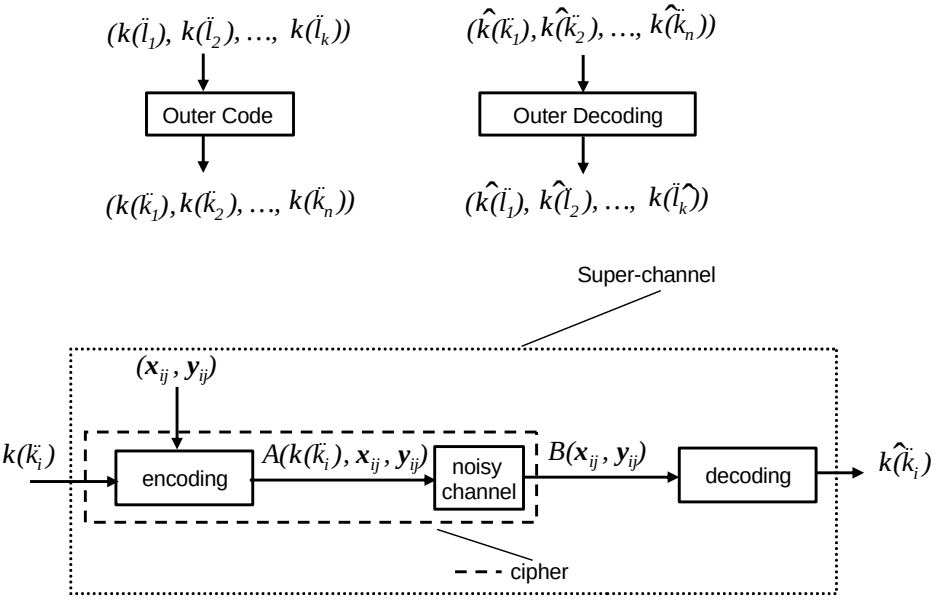

Figure 4: Related-key Attacks as Concatenated Codes

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 value[2]. We

[2]Note that the setting is different from that of the classical result on concatenated codes by [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

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 [19] 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, $\epsilon(N)$, reasonably small, in a single-key-recovery attack using N P/C pairs. Assuming that the superchannel is symmetric, let its capacity be $\mathcal{C}_S(N)$. (For small values of ϵ, $\mathcal{C}_S(N)$ is close to unity). 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 13 *Consider related-key-recovery attack Ψ 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. Denote the capacity of the superchannel by $\mathcal{C}_S(N)$. $\exists \psi$ such that*

$$\lim_{n,k\to\infty} \varepsilon(N,k,n,\psi) = 0 \text{ and } \nu(d,N,k,n) = \Lambda,\ \forall\ \Lambda \geq \frac{N}{d\mathcal{C}_S(N)}$$

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 [19, 4] says that $\exists$ an outer code ψ such that,

$$\lim_{n\to\infty} \varepsilon(N,k,n,\psi) = 0$$

for all constant $R \leq \mathcal{C}_S(N)$. As

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

this implies that

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

■

Theorem 13 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 14 *Given any $\alpha \in (0,1)$, and a single-key recovery attack Γ, $\exists$ N_α, a corresponding related-key recovery attack Ψ, and M_α, such that*

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

and

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

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.

Proof. Choose any $M_\alpha \geq \frac{d}{\mathcal{C}}$ and the corresponding ψ of Theorem 13. 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 13, $\exists\ n, k$ 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 15 *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 14 and Lemma 10,

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

■

In Theorem 13 and Corollaries 14 and 15 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 16) and to the cryptanalytic channel (Corollary 17) 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 16 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 17).

Corollary 16 *For related-key-recovery attack Ψ such that $\lim_{n\to\infty} \varepsilon(N, k, n, \psi) = 0$, and ν is constant, the minimum value of ν is $\overset{min}{N}\ \frac{N}{d\mathcal{C}_S(N)}$*

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, $\nu = \frac{nN_0}{kd} < \frac{N_0}{d\mathcal{C}_S(N_0)}$ and $\frac{n}{k} < \frac{1}{\mathcal{C}_S(N_0)}$, or the rate of the outer code, $\frac{k}{n}$, is larger than the capacity of the superchannel, $\mathcal{C}_S(N_0)$, contradicting the channel coding theorem. ■

Corollary 17 *For related-key-recovery attack* Ψ *with* $\nu(d, N, k, n) = constant$, *and*

$$\lim_{n\to\infty} \varepsilon(N, k, n, \psi) = 0$$

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

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

Proof. If not, Ψ 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. ■

Finally, the adversary need not maintain a constant value of ν, independent of n and k, for the attack. Using a result from [23], it can be shown that the tight bound of Corollary 16 is a tight asymptotic bound for attacks where ε is arbitrarily small, but ν is not required to be independent of k and n.

Corollary 18 *For related-key-recovery attack* Ψ *such that* $\lim_{n,k\to\infty} \varepsilon(N, k, n, \psi) = 0$, $\lim_{n,k\to\infty} \nu(N, d, k, n) \geq \frac{N}{d\mathcal{C}_S(N)}$

Proof. Follows from the fact that [23] $\lim_{n,k\to\infty} error = 0 \Rightarrow \lim_{n,k\to\infty} R \leq \mathcal{C}$. ■

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, (1) still holds. 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 Lemma 10 also holds.

The related-key attack uses the lists of estimates of each of the $\ddot{I}_i$, along with the value of the likelihood function, to rank combinations of estimates from the lists. This provides a list of estimates of $\ddot{\mathbf{I}}$, and Lemma 12 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 13), Corollaries 14 and 15 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 15 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.

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. This paper does not provide any experimental verification of the improvement in ν when the independent-key estimates are not unique.

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. We then 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 Reproduction of Matsui's original experiments on 8-round DES

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.

Figure 5 presents our results. The figure shows various graphs of the accuracy of the attack as a function of the number of solutions. That is, point (x, y) on a graph indicates that the correct key belongs to a list of the x best estimates with probability y. A specific graph corresponds to a specific value of N, the number of P/C pairs used. If N_8 denotes the number of P/C pairs required for a probability of error ϵ on 8-round DES (the graphs in Figure 5), and N_{16} that for the same error on full DES (the graphs in [15]), then:

$$N_8 = 1.49 \times 2^{-26} \times N_{16} \tag{10}$$

Thus the curve for 2^{43} P/C pairs in [15] corresponds to the one for 1.49×2^{17} (approximately, 200,000) P/C pairs in Figure 5.

As is clear from Figure 5, we have been able to reproduce Matsui's original attack with minor variations attributable to the minor differences in our experimental procedures. In the next section, we describe the calibration of the superchannel corresponding to the single-key recovery attack.

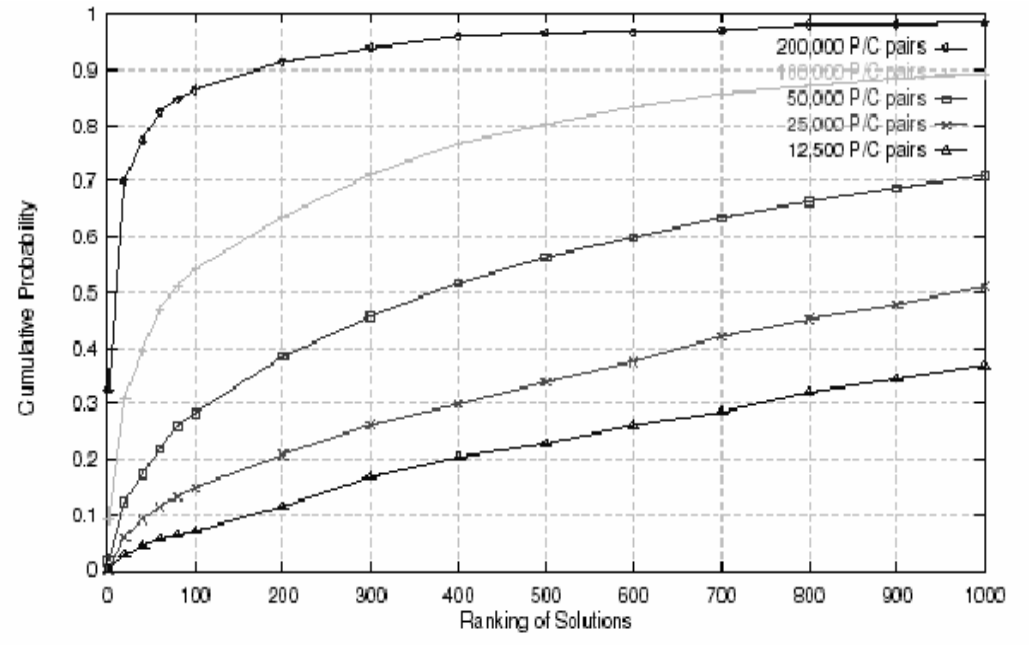

Figure 5: Accuracy of Key Recovery as a Function of List Size

5.3 Calibration of the Single-Key Recovery Attack of [16]

In Figure 6 we plot the probability of error, ϵ, for unique key-recovery, as a function of the number of P/C pairs, N. That is, the point $(1, y)$ on the graph corresponding to N P/C pairs in Figure 5 is the point $(N, 1-y)$ in Figure 6. In the CCM, Figure 6 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 $\frac{2(1.95\times 2^{-9})^2}{ln2}$), about $287,000$ P/C pairs, for any probability of error. In Matsui's original experiments, about $200,000$ P/C pairs resulted in a 40% error in unique-key recovery [15].

5.4 Limits on Related-Key Recovery

Recall from Theorem 13 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 7 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 16) 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 17).

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

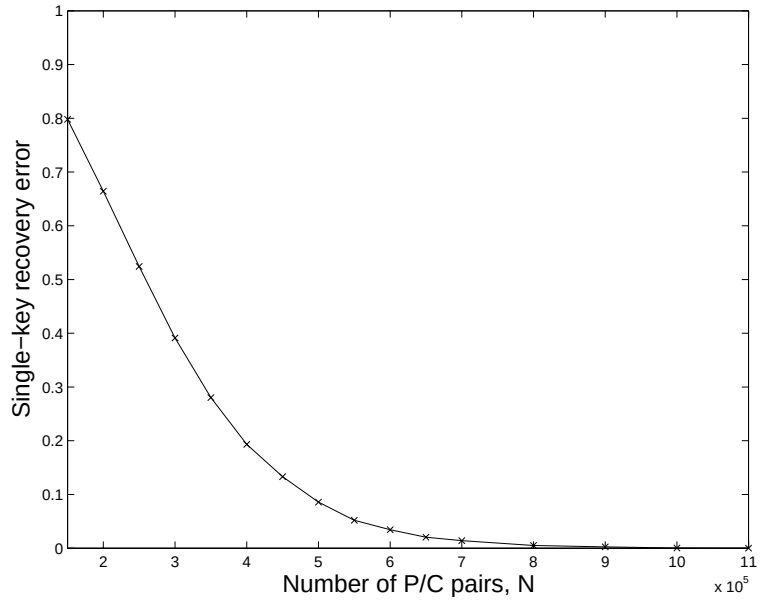

Figure 6: Single-Key Recovery Error, ϵ

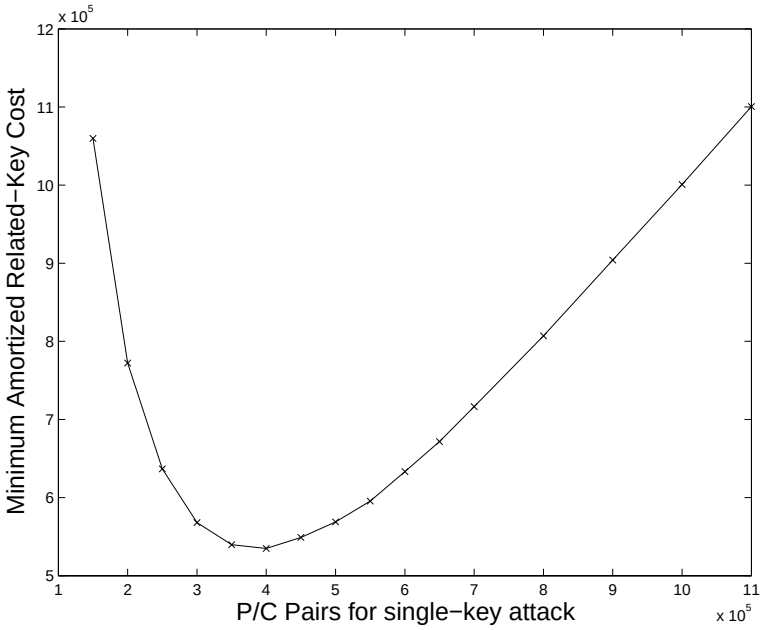

Figure 7: Minimum value of ν, $\frac{N}{d\mathcal{C}_S(\epsilon)}$

5.5 Related-Key Recovery on Reduced-Round DES

The experiment performed is as follows. k independent keys, ($\ddot{l}_1$, $\ddot{l}_2$, ..., $\ddot{l}_k$), are chosen. These are encoded using an (n, k) Reed Solomon code to obtain n related keys, ($\ddot{k}_1$, $\ddot{k}_2$, ..., $\ddot{k}_n$). The bits to be estimated by the attack are encoded independently of the other bits. That is, $\kappa(\ddot{l}_1)$, $\kappa(\ddot{l}_2)$, ..., $\kappa(\ddot{l}_k)$, are encoded to obtain $\kappa(\ddot{k}_1)$, $\kappa(\ddot{k}_2)$, ..., $\kappa(\ddot{k}_n)$, and the values of the other forty-four bits of each key are encoded separately. n single-key-recovery attacks are carried out using $\ddot{k}_1, \ddot{k}_2, ..., \ddot{k}_n$. The estimates obtained, $\widehat{\kappa(\ddot{k}_1)}$, $\widehat{\kappa(\ddot{k}_2)}$, ..., $\widehat{\kappa(\ddot{k}_n)}$ are RS-decoded to obtain $\widehat{\kappa(\ddot{l}_1)}$, $\widehat{\kappa(\ddot{l}_2)}$, ..., $\widehat{\kappa(\ddot{l}_k)}$. 1000 instances of the experiment are carried out for each value of (n, k). That is, for a given value of (n, k), the experiment is carried out for 1000 sets of keys $(\ddot{l}_1, \ddot{l}_2, ..., \ddot{l}_k)$, each set chosen uniformly at random. 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.

5.6 Experimental Results

In this section, we present the experimental results on 8-round DES, performed using 1000 instances of each related-key recovery attack. In Figure 8 we provide several plots. The solid-line plot corresponds to the single key recovery attack. 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 $1 - (1 - \varepsilon)^{\frac{1}{k}}$. 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.

The same results may be also be viewed as in Figure 9, which provides the values of ϵ beyond which it is more efficient to use the corresponding RS-encoded related-key attack, as a function of n. Notice that the plot is monotonic decreasing.

Finally, we observed that the related-key attack has an 11% lower value of ν than the independent-key attack, if $\varepsilon = 0.04$ and $k = 11$, a 22% lower value if $\varepsilon = 0.03$ and $k = 21$, and a 27% lower value if $\varepsilon = 0.17$ 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.

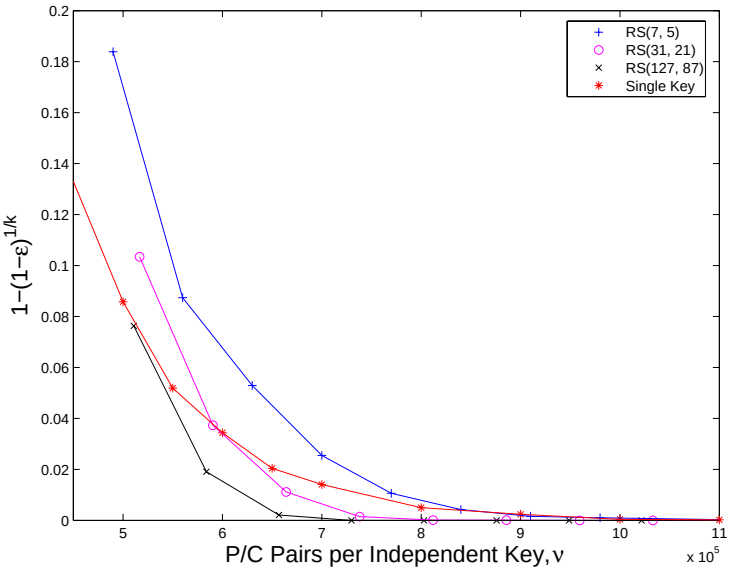

Figure 8: Equivalent single-key recovery error as a function of ν

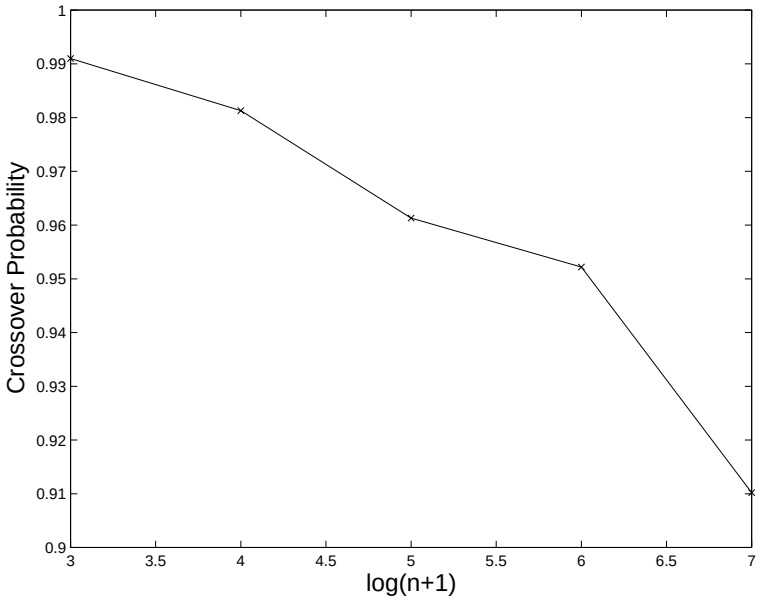

Figure 9: Crossover accuracy as a function of n

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 [11] 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.

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*, pages 2–21, London, UK, 1991. Springer-Verlag.

[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*, London, UK, 1995. Springer-Verlag.

[8] Howard Heys. A tutorial on linear and differential cryptanalysis. Technical Report CORR 2001-17, Centre for Applied Cryptographic Research, Department of Combinatorics and Optimization, University of Waterloo, March 2001.

[9] Thomas Jakobsen. Correlation attacks on block ciphers. Master's thesis, Dept. of Mathematics, Technical University of Denmark,, 1996.

[10] 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*, pages 212–222, London, UK, 1998. Springer-Verlag.

[11] Thomas Jakobsen. *Higher-Order Cryptanalysis of Block Ciphers*. PhD thesis, Dept. of Mathematics, Technical University of Denmark, 1999.

[12] 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*, London, UK, 1996. Springer-Verlag.

[13] 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*, pages 233–246, London, UK, 1997. Springer-Verlag.

[14] Lars R. Knudsen and David Wagner. Integral cryptanalysis. In *FSE '02: Revised Papers from the 9th International Workshop on Fast Software Encryption*, pages 112–127, London, UK, 2002. Springer-Verlag.

[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*, London, UK, 1994. Springer-Verlag.

[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*, pages 386–397, Secaucus, NJ, USA, 1994. Springer-Verlag New York, Inc.

[17] Darakhshan J. Mir. Related-key linear cryptanalysis of des. Master's thesis, School of Engineering and Applied Science, George Washington University, 2006.

[18] S. Murphy, F. Piper, M. Walker, and P. Wild. Maximum likelihood estimation for block cipher keys, 1994.

[19] Claude Shannon. A mathematical theory of communication. *Bell Systems Technical Journal*, 27:379–423, 1948.

[20] Harry L. Van Trees. *Detection, Estimation, and Modulation Theory, Part I*. John Wiley & Sons, 1968.

[21] S. Vaudenay. Decorrelation: A theory for block cipher security. *Journal of Cryptology: the journal of the International Association for Cryptologic Research*, 16(4):249–286, Sept. 2003.

[22] Serge Vaudenay. An experiment on des statistical cryptanalysis. In *CCS '96: Proceedings of the 3rd ACM conference on Computer and communications security*, pages 139–147, New York, NY, USA, 1996. ACM Press.

[23] 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. To appear.

[24] 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.

[25] David Wagner. Towards a unifying view of block cipher cryptanalysis. In *FSE '04: Eleventh International Workshop on Fast Software Encryption*, 2004.