Cryptanalysis of some recently-proposed multiple modes of operation

David Wagner

University of California, Berkeley
e-mail: daw@cs.berkeley.edu

Abstract. In a paper cryptanalyzing many triple modes of operation, Biham proposed four new triple modes and five new quadruple modes of operation for DES. It was conjectured that the complexity (in a particular threat model) of breaking the triple modes is at least 2^{112} and that the quadruple modes are more secure than any triple mode.
We present new attacks on all but one of the proposed modes. We can break all but two of Biham's proposed modes with at most 2^{56} off-line trial encryptions and between 2 and 2^{32} (depending upon the mode) chosen-IV chosen texts; another mode can be broken with somewhat more work. This raises questions about the suitability of the proposed modes, and provides further evidence for the fragility of inner chaining; however, we emphasize that our results do not disprove Biham's conjectures, as we rely on an extended attack model which admits more powerful adversaries who can mount chosen-IV queries, a capability denied to them in Biham's model.

1 Introduction

DES is the most thoroughly-analyzed cipher in the open literature, but after more than two decades, it is reaching the end of its useful lifetime: the DES 56-bit key-length is simply too short to be secure against serious keysearch efforts. Therefore, there is great interest in the search for a multiple mode of operation for DES which provides increased strength against exhaustive keysearch while retaining the high level of analysis and confidence that single-DES currently offers.

Biham [Bih96] analyzed a great many triple modes of operation, and broke every mode considered except the commonly-used triple-DES-ECB mode (when used with some outer chaining technique). Unfortunately, due to its short 64-bit block length, triple-DES-ECB has some shortcomings: it is susceptible to dictionary attacks (when 2^{64} known texts are available) and matching-ciphertext attacks (where partial information about the plaintext is recovered by using the birthday paradox, when 2^{32} known texts are available).

To improve this state of affairs, Biham proposed 9 new block modes and 2 new stream modes of operation for DES. The complexity of attacking these new modes is conjectured to be at least 2^{112}. The quadruple modes were conjectured to be more secure than any triple mode; furthermore, the complexity of attacking two of the quadruple modes was conjectured to be at least 2^{128}.

This paper shows that, when we allow chosen-IV chosen-text attacks, most of the proposed modes are not significantly more secure than single-DES. We provide new attacks against all but one of the modes.

Note that Biham's studies were premised on a more restrictive threat model that did not admit chosen-IV attacks, so our results do not disprove Biham's conjectures; but our position is that these new results raise questions about the security of Biham's proposed modes and illustrate the application of general techniques for cryptanalysis of multiple modes of operation. See Section 3 for more discussion on this point.

The paper is organized as follows. Section 2 establishes some notation and other background, and Section 3 discusses our threat model. Section 4 shows how to attack two important classes of modes using a divide-and-conquer strategy, and applies this result to attack six of Biham's proposed modes. Section 5 shows how to attack four more of Biham's modes using narrow-pipe attacks. Finally, Section 6 discusses some implications of our results, and Section 7 wraps up the paper with some concluding remarks.

2 Preliminaries

Biham developed a concise notation for multiple modes which is worth summarizing here. All of his new modes are derived from the standard DES modes of operation—ECB, CBC, CFB, and OFB—as well as their corresponding decryption modes—ECB^{-1}, etc. The notation CBC|CFB refers to the mode where the output of DES-CBC encryption is fed to the input of DES-CFB encryption; the | operator can be extended to triple and higher-order modes. The notation OFB[CBC] refers to a mode which applies OFB to its input, then encrypts with CBC mode, and finally applies the same OFB keystream to that result. (Note: the streams xored into the input and the output of CBC are generated from a single DES key, and therefore are *the same*.) This can be generalized to modes such as OFB[CBC,CFB], where we apply OFB, then CBC, then OFB again, then CFB, and then OFB once more. (Again, all the OFB output streams are the same!) The notation OFB→CBC refers to a stream mode which applies CBC encryption to the keystream generated by OFB mode, and xors the result to the plaintext. We can of course use the → operator to define triple and higher-order modes, too.

For clarity, we will attempt to use the same notation for plaintext, ciphertext, etc. throughout this note. We write $P_0, P_1, \ldots$ (respectively $C_0, C_1, \ldots$) for the blocks of the plaintext (resp. ciphertext). We let $K0, K1, \ldots$ denote the 56-bit DES keys, and write $IV0, IV1, \ldots$ for the corresponding IVs. We number the keys $K0, K1, \ldots$ according to the order that the single-mode appears in this notation: for instance, in OFB[CBC,CBC^{-1}], the OFB-mode is keyed with $K0$, the CBC with $K1$, and the CBC^{-1} with $K2$. When multiple plaintext/ciphertext pairs are obtained in an attack run, we write $P[j]$ for the full plaintext of the

j-th message, write $P_i[j]$ for the i-th block of $P[j]$, and so on. We let $E_k(x)$ stand for the single-DES encryption of the input block x under the key k.

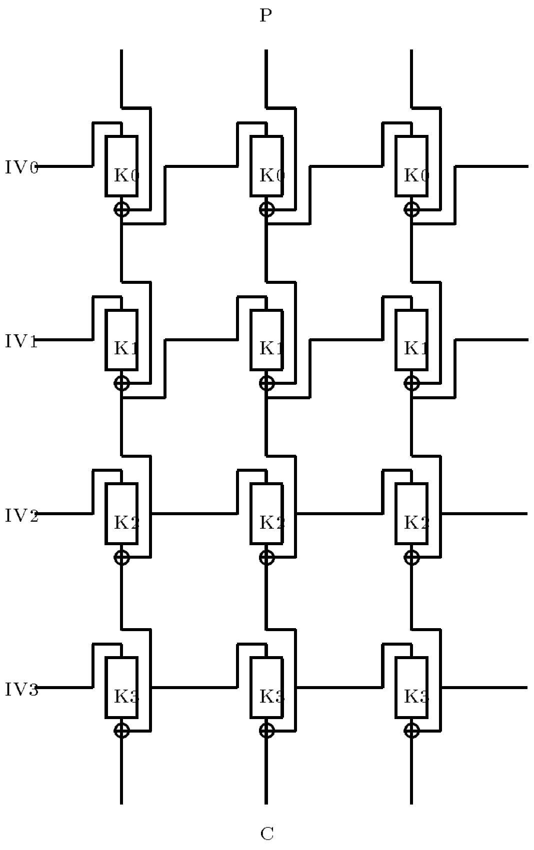

Fig. 1. The CFB|CFB|CFB^{-1} |CFB^{-1} mode.

As an example of this notation, we depict the the CFB|CFB|CFB^{-1} |CFB^{-1} mode in Figure 1.

All of our attacks recover the secret DES keys. The basic ideas behind the attacks are not entirely novel; many of them are applications of the general tools worked out by Coppersmith, Johnson, and Matyas [CJM97] and Biham [Bih94a,Bih94b,Bih96].

The nine new block modes which Biham proposed are

1. OFB[CBC,CBC^{-1}],
2. OFB[CFB,CFB^{-1}],
3. OFB[CBC,CBC],
4. OFB[CFB,CFB],

5. CBC|CBC|CBC^{-1} |CBC^{-1},
6. CFB|CFB|CFB^{-1} |CFB^{-1},
7. OFB[CBC,CBC,CBC^{-1}],
8. OFB[CBC,CBC,CBC], and
9. OFB[CFB,CFB,CFB].

The proposed stream modes are

1. OFB→CBC→CBC, and
2. OFB→CFB→CFB.

In the following sections, we find new attacks on all of these except OFB[CBC, CBC^{-1}].

3 Security model

In this section, we examine the attack model. The opponent is assumed to have the necessary computational power to perform 2^{56} off-line trial encryptions. We assume (as is standard) that the adversary can perform known-plaintext, chosen-plaintext, and chosen-ciphertext attacks.

So far we have not deviated from Biham's model. We list below three important differences.

3.1 Controlling IVs

The most important difference between the two models comes when we examine the treatment of IVs.

In our model, a mode is essentially a mini-protocol specifying how to perform secure message transport. To send the secret message P, one enciphers P under the appropriate multiple mode with key K and with randomly chosen IVs $IV0, \ldots, IVn$, transmitting the bundle $IV0, \ldots, IVn, C$ over the insecure medium; the receiver decrypts C with the specified IVs under the shared key K and recovers the decrypted message P'. The subtlety comes when we introduce active attackers with the ability to perform chosen-ciphertext attacks: such adversaries are free to specify any ciphertext C along with any set of IVs they wish, and they will receive the decryption P' of that ciphertext.

Our attack model captures this notion. Because we allow chosen-ciphertext attacks, we also allow (as a natural consequence) chosen-IV chosen-ciphertext attacks. It is worth noting that this choice induces a slight assymetry between chosen-plaintext and chosen-ciphertext attacks: adversaries may control the IV in chosen-ciphertext attacks, but not in chosen-plaintext attacks.

In contrast, Biham did not consider chosen-IV attacks; even known-IV attacks were mentioned only in a few special cases. His model is more elegant and cleaner for analysis; for instance, the symmetry ensures that the security factor for a mode is the same as for its inverse. Also, attacks are all the more compelling when they are performed in Biham's more restrictive model. Finally, Biham's attacks remain applicable even when special measures to protect the IV are taken, whereas our attacks may be stopped by such measures.

We take the conservative philosophy that our model should allow adversaries considerable leeway; if the cryptosystem can stand up to attack in such a model, our assurance of security will be all the greater. Part of our justification for this approach is that triple DES with outer chaining already offers pretty good security, with only a few shortcomings: if we want to do better, our threshold should be quite high.

Our attacks will take advantage of this ability to control the IV, so they are not directly comparable to Biham's results. However, a number of the chosen-IV attacks can be converted to known-IV attacks with only a minor increase in the complexity of cryptanalysis, so some comparisons may be possible. See Section 6. Even where we are not aware of known-IV attacks, we view our chosen-IV attacks as certificational weaknesses that should at the very least raise warning flags about the security of the modes in question.

This subject is not yet exhausted. See Section 6 for some simple countermeasures to resist chosen-IV attacks, some counter-countermeasures, and their implications for the interpretation of our results.

3.2 Adaptive attacks

Biham's model also differs from ours in another respect: we allow adaptive chosen-text attacks, whereas Biham did not consider adaptive attacks. Moreover, Biham generally required only one encrypted stream for his analyses. In contrast, all of our attacks are cast in the language of adaptive attacks.

We view this distinction as relatively minor. All of our adaptive attacks can be easily converted to non-adaptive attacks with negligible increase in complexity (and, occasionally, a substantial increase in the number of messy details); in short, the adaptivity is merely convenient, not fundamental.

3.3 The cost of chosen texts

One philosophical point is that we try to be explicit about the resource requirements of our attacks, listing separately the number of chosen texts, offline decryptions, and memory words needed. The reader is then free to assign appropriate costs to each resource, according to his or her security environment.

It would be simpler to label each attack with a simple complexity measure that equates the cost of one chosen text with the cost of one trial decryption. Indeed, such a measure has great benefits for simplifying analysis, summarizing results, and comparing modes; and it is a very useful first approximation. The drawback is that highly theoretical attacks needing 2^{56} chosen texts may be equated with more serious attacks needing only 2^{56} trial decryptions. In practice, that distinction can be critical. Therefore, where possible, we aim to improve the quality of the approximation by using more explicit complexity measures.

4 Divide-and-conquer attacks

First, we list some fairly elementary attacks on several modes. These all have the flavor of "divide-and-conquer" algorithms: namely, we isolate the effect of each subkey with a chosen-ciphertext probe, and then recover each subkey with a 2^{56} exhaustive keysearch.

By the end of this section, we will see how to attack the block modes OFB[M1,M2,...,Mn] and M0|M1|M2|...|Mn, for any n, in the special case where each mode Mj is either CFB or CFB^{-1}. The intuition is that, in such modes, we have the relation

$$C_0 = P_0 \oplus E_{K0}(IV0) \oplus \ldots \oplus E_{Kn}(IVn) \tag{1}$$

on the first block; this is highly linear and, therefore, highly suspicious.

We will also see how to attack stream modes of the form OFB→M1→ ... →Mn, if each mode Mj is one of OFB, CBC, CBC^{-1}, CFB, CFB^{-1}. The idea is that we can apply a divide-and-conquer attack that isolates the effect of the last key Kn (with a single chosen-ciphertext query that probes IVn); we then strip off the last mode and continue iteratively.

4.1 OFB[CFB,CFB]

Our attack on the OFB[CFB,CFB] mode is composed of three phases; each phase isolates the effect of one DES key Kj. First, we recover the key $K1$ used in the first CFB mode by using one chosen-IV chosen-ciphertext query: namely, we isolate the effect of $K1$ by probing $IV1$. In the second phase, we recover the key $K2$ by probing $IV2$ with a similar chosen-ciphertext query. Finally, $K0$ is recovered by exhaustive keysearch.

In the first phase, we probe $IV1$ to isolate the effect of $K1$, and recover $K1$ with a 2^{56} exhaustive keysearch. Let $P[0], C[0]$ be a known plaintext/ciphertext pair with known IVs. We construct a chosen ciphertext query $C[1]$ as follows. Pick $IV1[1] \neq IV1[0]$, set $IV0[1] = IV0[0], IV2[1] = IV2[0]$, take $C[1] = C[0]$, and obtain the decryption $P[1]$ of the new ciphertext. Note that, by Equation 1,

$$P_0[0] \oplus P_0[1] = E_{K1}(IV1[0]) \oplus E_{K1}(IV1[1]).$$

Therefore we may find $K1$ by a 2^{56} exhaustive keysearch, recognizing the right key value when the above equation holds; with high probability, we expect no wrong key value to survive the check.

The second phase recovers $K2$ in an entirely analogous fashion, this time probing $IV2$ instead of $IV1$.

Finally, in the third phase we perform a 2^{56} exhaustive search over $K0$ (the only remaining unknown key value). Therefore, the total complexity of the attack is two chosen-ciphertexts and $5 \cdot 2^{56}$ off-line trial encryptions.

4.2 OFB[CFB,CFB^{-1}]

The OFB[CFB,CFB^{-1}] mode can be broken in a way entirely analogous to the cryptanalysis of OFB[CFB,CFB]: probe $IV1$ in one chosen-ciphertext query to recover $K1$, then probe $IV2$ to learn $K2$, and exhaustively search over $K0$. So the OFB[CFB,CFB^{-1}] mode, too, can be broken with two chosen-ciphertexts and $5 \cdot 2^{56}$ off-line trial encryptions.

4.3 OFB[CFB,CFB,CFB]

The OFB[CFB,CFB,CFB] mode can also be broken with the same technique. For this mode, we need three chosen-ciphertexts and $7 \cdot 2^{56}$ off-line trial encryptions.

4.4 CFB|CFB|CFB^{-1} |CFB^{-1}

This mode is also easy to break using the same techniques. (Note that the CFB|CFB|CFB^{-1} |CFB^{-1} mode is illustrated in Figure 1.) As before, in the first phase we can probe $IV0$ to isolate the effect of $K0$ and recover $K0$ by exhaustive search; continue to recover the rest of the keys. In this way, we can break the CFB|CFB|CFB^{-1} |CFB^{-1} mode with a total of three chosen-ciphertexts and $7 \cdot 2^{56}$ off-line trial encryptions.

4.5 OFB→CBC→CBC

OFB→CBC→CBC mode is characterized by the relation

$$C_0 = P_0 \oplus E_{K2}(IV2 \oplus E_{K1}(IV1 \oplus E_{K0}(IV0))).$$

In the first phase of our attack, we probe $IV2$ to isolate the effect of $K2$. More precisely, let $P[0], C[0]$ be a known plaintext/ciphertext pair with IVs $IV0[0], IV1[0], IV2[0]$, and construct a chosen-ciphertext query as follows. Set $C[1] = C[0]$, pick $IV2[1] \neq IV2[0]$, and set $IV0[1] = IV0[0], IV1[1] = IV1[0]$. Next issue a chosen-ciphertext query for the $C[1], IVj[1]$ to get $P[1]$. Finally note that

$$IV2[0] \oplus IV2[1] = D_{K2}(P_0[0] \oplus C_0[0]) \oplus D_{K2}(P_0[1] \oplus C_0[1]);$$

this relation lets us recover $K2$ with a 2^{56} exhaustive search.

The second phase of the attack probes $IV1$ in a similar way to recover $K1$. Finally, $K0$ can be obtained in a third phase by brute force. In sum, this cryptanalysis requires $5 \cdot 2^{56}$ off-line trial encryptions and two chosen-ciphertexts.

4.6 OFB→CFB→CFB

Our attack on OFB→CFB→CFB mode proceeds in a very similar way to that described in the previous paragraph. We probe $IV2$ in a chosen-ciphertext attack, which allows us to isolate the effect of $K2$ by the following relation:

$$C_0[0] \oplus C_0[1] \oplus P_0[0] \oplus P_1[1] = E_{K2}(IV2[0]) \oplus E_{K2}(IV2[1]).$$

Then $K1$ is recovered analogously, and $K0$ by exhaustive keysearch. The total complexity of this attack is two chosen-ciphertext queries and $5 \cdot 2^{56}$ off-line trial encryptions.

5 Narrow-pipe attacks

In this section, we describe a number of narrow-pipe attacks. (By a "narrow pipe", we mean a data channel that is relatively narrow—only 64 bits wide, for instance.) The basic technique is to identify some narrow pipe through which all diffusion is channeled; then you generate a bunch of texts and look for a collision in that narrow pipe. The birthday paradox assures us that we will find a collision in the narrow pipe relatively quickly (within $2^{n/2}$ texts, for a n-bit pipe). Then we hope (1) that we can recognize the collision by looking only at the plaintext and ciphertext, and (2) that we can use that knowledge to deduce some relation which isolates the effect of just one DES key. When the attack is designed correctly, we will be able to find recognizable collisions in the narrow pipe that let us deduce important information about some key Ki standing alone. After recovering Ki with a 2^{56} exhaustive keysearch, we remove the effect of that key and attempt to solve the reduced mode by iterating the attack.

In this section, we show how to break the CBC|CBC|CBC^{-1} |CBC^{-1} block mode, as well as the OFB[CBC,CBC], OFB[CBC,CBC,CBC], and OFB[CBC, CBC,CBC^{-1}] modes.

5.1 CBC|CBC|CBC^{-1} |CBC^{-1}

To break CBC|CBC|CBC^{-1} |CBC^{-1} (see Figure 2), we first recover $K0$ by probing $IV1$. Let $P[0], C[0]$ be a known plaintext/ciphertext pair with known IVs, and build a chosen ciphertext query as follows. Pick $IV1[1] \neq IV1[0]$, set $IVj[1] = IVj[0]$ for $j \neq 1$, take $C[1] = C[0]$, and obtain the decryption $P[1]$ of the new ciphertext. Note that

$$IV1[0] \oplus IV1[1] = E_{K0}(IV0[0] \oplus P_0[0]) \oplus E_{K0}(IV0[1] \oplus P_0[1]).$$

Therefore we may find $K0$ by a 2^{56} exhaustive keysearch, recognizing the right key value when the above equation holds; with high probability, the check will eliminate all incorrect guesses at the key.

Once we've learned $K0$ with 2^{56} work and one chosen-ciphertext query, we can peel off the effect of $K0$ and reduce the problem to that of breaking the CBC|CBC^{-1} |CBC^{-1} mode.

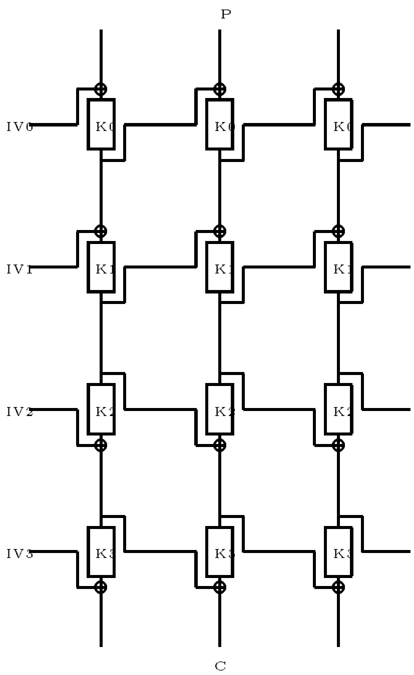

Fig. 2. The CBC|CBC|CBC^{-1} |CBC^{-1} mode.

Biham conjectured that the quad mode CBC|CBC|CBC^{-1} |CBC^{-1} is more secure than any triple mode. Note that our present attack does not immediately disprove Biham's conjecture, since our reduction relies on mounting a chosen-IV chosen-ciphertext query, which is not allowed in Biham's security model. In Section 6.2 we extend it to work with only known-IV queries, which brings us a step closer to Biham's model.

Finishing the attack. We now describe how to finish the attack on CBC| CBC|CBC^{-1} |CBC^{-1}. All that remains is to analyze the triple mode CBC| CBC^{-1} |CBC^{-1}. Biham has showed how to break this triple mode with 2^{68} chosen plaintexts and 2^{66} work [Bih96]. Nonetheless, in our security model, 2^{68} texts are quite a high barrier, and one might wonder whether there are more efficient attacks.

The answer is yes. We present next a new attack on CBC|CBC^{-1} |CBC^{-1} which requires only 2^{32} chosen-ciphertext chosen-IV queries and $6 \cdot 2^{56}$ trial encryptions. This can be used as a subroutine to develop a full attack on the CBC|CBC|CBC^{-1} |CBC^{-1} quad mode with roughly equivalent complexity.

Breaking CBC|CBC^{-1} |CBC^{-1}. We recover $K2$ by probing $IV1$. Fix arbitrary IVs $IV0, IV2$. We construct 2^{32} chosen-ciphertext queries, as follows. For each i, pick $C_0[i]$ and $IV1[i]$ randomly, and let $C_1[i] = IV1[i]$. Now we obtain the decryptions $P[i]$ of those 2^{32} ciphertexts. We search for i, j such that $P[i] = P[j]$ (using a hash table, so as to avoid increasing the complexity of the attack). Note that two plaintexts will certainly match in the first block (i.e. $P_0[i] = P_0[j]$) if

$$IV1[i] \oplus IV1[j] = E_{K2}(C_0[i] \oplus IV2) \oplus E_{K2}(C_0[j] \oplus IV2) \tag{2}$$

holds, because then the resulting collision between the second and third DES layers will necessarily propagate up to the plaintext. Moreover, if Equation 2 holds, then in fact the two plaintexts $P[i], P[j]$ will match in their entirety. (This is because the value at the bottom of the third DES layer at the second block is $C_1 \oplus E_{K2}(C_0 \oplus IV2)$; now the choice of $C_1[i], C_1[j]$ ensures that $C_1[i] \oplus E_{K2}(C_0[i] \oplus IV2) = IV1[i] \oplus (E_{K2}(C_0[j] \oplus IV2) \oplus IV1[i] \oplus IV1[j]) = IV1[j] \oplus E_{K2}(C_0[j] \oplus IV2) = C_2[j] \oplus E_{K2}(C_0[j] \oplus IV2)$, so we will get a collision at the bottom of the third layer, and this will necessarily propagate up to the plaintext.) Since we generated 2^{32} ciphertexts and the block size is 64 bits, by the birthday paradox with high probability we will find one pair i, j satisfying Equation 2, and so with high probability we will see $P[i] = P[j]$ for some i, j. On the other hand, because P is two blocks (128 bits) long, the chances of seeing a chance match $P[i'] = P[j']$ by accident is very low. Therefore, we expect to see one match $P[i] = P[j]$, and we can conclude that for such i, j Equation 2 must hold.

Once we've found a pair i, j where Equation 2 holds, we can use it to isolate the effect of $K2$. This lets us recover $K2$ using a 2^{56} exhaustive keysearch. Finally, knowing $K2$ lets us reduce the problem to that of breaking CBC|CBC^{-1} mode, which can be done by standard techniques without any increase in complexity. This lets us break the triple mode CBC|CBC^{-1} |CBC^{-1} with $6 \cdot 2^{56}$ trial encryptions and 2^{32} chosen-ciphertext chosen-IV queries. (In fact, the attack can be extended to work just as efficiently with known-IV chosen-ciphertext queries instead of chosen-IV queries; it just becomes a bit messier to describe.)

Pulling it all together. To summarize, we can apply these techniques to break the CBC|CBC|CBC^{-1} |CBC^{-1} quad block mode with 2^{32} chosen-ciphertext queries and $7 \cdot 2^{56}$ trial encryptions.

Note that we could dramatically reduce the number of chosen-ciphertext queries needed if there were a better way to break the triple mode CBC|CBC^{-1} | CBC^{-1}.

5.2 OFB[CBC,CBC]

For the OFB[CBC,CBC] mode, we use another narrow-pipe attack combined with a birthday argument to find a collision in the OFB streams generated by two different messages. Generate 2^{32} chosen-ciphertext queries, as follows. Fix a 64-bit constant c, fix $IV1, IV2$, and let $C[i] = (c, c, c, c, c)$ for all i. The only value that varies will be $IV0[i]$, which we pick randomly. Obtain the decryptions $P[i]$ of those chosen ciphertexts. Now we search for i, j such that $E_{K0}(IV0[i]) = IV0[j]$; this relation ensures that the two OFB streams for $P[i], P[j]$ will match up (except that they will be out of phase by one block). Of course, given such an i, j, we can recover $K0$ with a 2^{56} exhaustive keysearch.

How can we recognize such a fortunate event? Note that peeling off the second CBC mode during decryption of $C[i]$ leaves $(?, ?, d, e, f)$, while for $C[j]$ we get $(?, d, e, f, ?)$. Furthermore, peeling off the first CBC mode we get $(?, ?, ?, g, h)$ for $C[i]$ and $(?, ?, g, h, ?)$ for $C[j]$. (A word on notation: the question marks "?" just represent arbitrary unknown values; two entries both marked with a "?" need not be equal.) In other words, we can recognize i, j such that $E_{K0}(IV0[i]) = IV0[j]$ by the requirement that $P_3[i] = P_2[j]$ and $P_4[i] = P_3[j]$; false alarms should be very rare, and with 2^{32} chosen ciphertexts, the birthday paradox reassures us that we expect to find at least one such i, j.

Once we've found i, j such that $E_{K0}(IV0[i]) = IV0[j]$, we can recover $K0$ with a 2^{56} exhaustive search. This reduces the problem of breaking OFB[CBC, CBC] to that of breaking (a mode very similar to) CBC|CBC. The latter mode can be broken efficiently with standard techniques, and in fact we can recover $K1, K2$ with one chosen ciphertext and $3 \cdot 2^{56}$ trial encryptions. In total, we can break OFB[CBC,CBC] with 2^{32} chosen ciphertexts and $4 \cdot 2^{56}$ trial encryptions.

5.3 OFB[CBC,CBC,CBC]

OFB[CBC,CBC,CBC] can be broken similarly. First, we recover $K0$ using the same technique as described above for OFB[CBC,CBC]. Then all we need to do is break (a mode very similar to) CBC|CBC|CBC, and that can be done with two chosen ciphertext queries and $5 \cdot 2^{56}$ trial encryptions. In total, our cryptanalysis of OFB[CBC,CBC,CBC] needs 2^{32} chosen ciphertexts and $6 \cdot 2^{56}$ trial encryptions.

5.4 OFB[CBC,CBC^{-1}]

We are not aware of any strong attacks on the OFB[CBC,CBC^{-1}] mode.

We have a narrow-pipe attack that recovers the key with 2^{112} offline trial decryptions, 2^{32} chosen ciphertexts, and no memory; but because this result is so weak, we will refrain from describing it here.

We also have an attack that requires 2^{66} known-IV chosen-ciphertext queries, 2^{56} offline trial decryptions, and 2^{66} memory. This too is highly unrealistic, but we will sketch the attack here for completeness. We proceed much as in Section 5.2, with the additional complication that we must also force an internal feedback channel to match. Choose 2^{64} ciphertexts $C[i] = (c, c, c, c, c)$ for all i. We seek i, j such that $E_{K0}(IV0[i]) = IV0[j]$, which will ensure that the two OFB streams match up (out of phase by one block); we also require that $E_{K2}(C_0[i] \oplus E_{K0}(IV0[i]) \oplus IV2[i]) = IV2[j]$, which yields a collision (out of phase by one block) in the CBC^{-1} layer's internal feedback channel. These two conditions ensure that we can recognize such a pair i, j by the condition $P_{2...4}[i] = P_{1...3}[j]$. Furthermore, the birthday paradox predicts that we will encounter one such i, j; once we've recognized it, we can use the known IVs to recover $K0$ with a 2^{56} exhaustive keysearch. Then the rest of the key material can be obtained with a meet-in-the-middle search.

In short, we are unable to make much progress on the analysis of OFB[CBC, CBC^{-1}], and so we leave it as an open question for others to examine.

5.5 OFB[CBC,CBC,CBC^{-1}]

We present an attack that breaks the OFB[CBC,CBC,CBC^{-1}] mode with one chosen ciphertext query and 2^{112} work. This is an unrealistic attack, but it shows that this quadruple mode does not attain the strength one might ideally hope for in a quad mode, when chosen-IV queries are a threat.

First, we probe $IV2$ to isolate the effect of $K0, K1$. Let $P[0], C[0]$ be a known plaintext/ciphertext pair with known IVs $IVi[0]$. We construct a chosen ciphertext query $C[1], IVi[1]$ by taking $C[1] = C[0]$, letting $IVi[1] = IVi[0]$ for $i = 0, 1, 3$, and picking arbitrary $IV2[1]$ different from $IV2[0]$. Then we have the relation

$$IV2[0] \oplus IV2[1] = E_{K1}(P_0[0] \oplus IV1 \oplus E_{K0}(IV0)) \oplus E_{K1}(P_0[1] \oplus IV1 \oplus E_{K0}(IV0)),$$

which allows us to isolate the effect of $K0, K1$. Now we recover $K0, K1$ with a 2^{112} exhaustive keysearch.

Finally, once we've learned $K0, K1$, we can recover $K2, K3$ with a second exhaustive keysearch. (In fact, we could use the meet-in-the-middle attack on double-DES to recover $K2, K3$, but this will not reduce the total complexity of the full attack significantly.)

The total complexity of the attack is $3 \cdot 2^{112}$ offline trial encryptions and one chosen-ciphertext query. This shows that the quadruple mode OFB[CBC,CBC, CBC^{-1}] is no stronger than triple-DES-ECB (with outer chaining) against chosen-IV chosen-text key-recovery attacks, and so the fourth DES layer seems wasted.

It is interesting to note that the present attack does not apply to the triple mode OFB[CBC,CBC^{-1}], even though this might seem like a paradox at first glance. This leaves open the counter-intuitive possibility that the OFB[CBC, CBC^{-1}] triple mode might well be stronger than the OFB[CBC,CBC,CBC^{-1}] quadruple mode.

6 Discussion

All of our attacks have relied on the ability to control the IV afforded by our attack model. This raises the issue of whether it is possible to prevent these attacks with simple countermeasures. The answer seems to be mixed—yes, there are some simple countermeasures, but they have limitations. We survey some possible approaches here.

6.1 Countermeasures

Encrypt the IVs. One suggestion is to encrypt the IVs before transmission (with, say, triple-DES or quadruple-DES), rather than sending the IVs in the clear. This thwarts the attackers ability to choose the exact values of the IVs.

On the other hand, this approach has a weakness, when one encrypts each IV independently: the attacker can still reuse old IV values in new chosen ciphertexts. Some of our attacks (suitably modified) can be converted to work against this protocol, when (1) they only rely on the ability to force IVj (for some choices of j) to be the same for all chosen ciphertexts and (2) the actual value of IVj is irrelevant. As an illustration, we show that OFB→CBC→CBC is no safer with this protocol than before: take

$$IV2[0] = IV2[2] = a,\ \ IV2[1] = IV2[3] = b,$$
$$IV1[0] = IV1[1] = c,\ \ IV1[2] = IV1[3] = d$$

for some unknown a, b, c, d, and force IV0 to be constant; then we have the identity

$$D_{K2}(P_0[0] \oplus C_0[0]) \oplus D_{K2}(P_0[1] \oplus C_0[1]) = D_{K2}(P_0[2] \oplus C_0[2]) \oplus D_{K2}(P_0[3] \oplus C_0[3]),$$

which lets us recover $K2$ with 2^{56} work and four chosen-ciphertext queries, and $K1$ will fall soon thereafter. So encrypting the IVs is no guarantee of safety.

MAC the IVs. Another natural reaction is to simply insist that senders apply a MAC to the IVs which receivers must verify before decrypting. By protecting the integrity of the IVs, this stops chosen-IV attacks.

This approach still leaves the users open to known-IV attacks, when they exist. Some modes are susceptible to known-IV attacks; others may not be. See below for a few illustrations of this danger. In general, the known-IV attacks that we know of usually require more texts than their chosen-IV counterparts, so MACing the IVs may reduce the threat level.

Another argument against this approach is based on engineering considerations. Now we have a new protocol, which is more complicated, and which introduces an whole new primitive to the mix. We end up replacing one failure mode with two: if either the encryption algorithm or the MAC is compromised, then the message keys must be recovered. It is perhaps imprudent to rely on the security of the MAC to protect confidentiality: just as conservative cryptographic design calls for independent session keys for authentication and confidentiality algorithms (to limit the impact of the compromise of any one algorithm), we would do well to avoid linking the security of our MAC with the security of our encryption algorithm.

We could perhaps first encrypt and then authenticate the IVs, to stop both known-IV attacks and attacks which attempt to replay old IVs. However, adding this much complexity to the system may begin to test the limits of one's comfort zone; at the least, more analysis seems needed.

Add redundancy to IVs. Coppersmith et al. [CJM97] have applied a novel countermeasure to stop a chosen-IV attack we discovered on their original CBCM proposal. They limit the possible values for each IV to a small subset: one IV is fixed at 0, and the other 64-bit IV has 44 of its bits fixed at 0. This redundancy limits the ability of an attacker to control the IV, and counters the attack we found.

This is a very clever trick, but it only seems useful in certain cases. Adding redundancy to IVs will not stop most of the attacks listed in this paper. Fixing certain IVs at 0 would deter many of the attacks, but it seems that such a measure could adversely affect security in other ways for a number of the modes proposed by Biham. It seems possible that this countermeasure may introduce as many problems as it solves, and so we are wary of depending upon it for security.

General notes. This is by no means an exhaustive list of available remedies. Nonetheless, we can make some comments that seem broadly applicable.

Many of the obvious countermeasures have not yet been subjected to concerted analysis, and we have attempted to show that there are some pitfalls to watch out for.

Still, it seems likely that techniques can be developed to protect the IVs quite thoroughly, for some (if not all) of Biham's modes. Of course, one has to use them, and use them with extreme care; it is details like this that plague real implementations. The central question is this: will such countermeasures prove cost-effective? or will these advanced modes suffocate under the weight of the extra precautions they require? More research is needed.

6.2 Extending our results to other security models

Stopping chosen-IV attacks is not enough, if the basic ideas behind those attacks can be leveraged into a sharper attack. To illustrate the point, we note that a number of our chosen-IV attacks can be converted to known-IV attacks. Usually, this increases the number of texts needed to mount the attack. These known-IV attacks are invariably more difficult to describe, and—perhaps—more difficult to discover, than their chosen-IV counterparts.

For example, all of Section 4's divide-and-conquer attacks on triple modes can be modified to work with 2^{64} known-IV chosen plaintexts. Use a birthday attack to find two texts with matching values of $IV0, IV1$; then that pair lets you probe $IV2$ and thus recover $K2$, and $K1$ is recovered similarly.

The number of known-IV queries can be reduced by using meet-in-the-middle techniques. For instance, one can break OFB[CFB,CFB] with 2^{32} known-IV chosen plaintexts. Use a birthday attack to find two texts $C[i], C[j]$ with $IV0[i] = IV0[j]$. We learn that

$$E_{K1}(IV1[i]) \oplus E_{K1}(IV1[j]) = E_{K2}(IV2[i]) \oplus E_{K2}(IV2[j]),$$

which lets us recover $K1, K2$ with complexity 2^{56} by a standard meet-in-the-middle attack. (The straightforward implementation of that attack also requires 2^{56} space, though the space requirements can be dramatically reduced by using parallel collision search algorithms [OW96].)

Applying these techniques, we can convert our chosen-IV attacks to attacks which need 2^{32} known-IV chosen texts and $O(2^{56})$ work, for all the triple modes in Section 4, as well as for OFB[CBC,CBC]. OFB[CFB,CFB,CFB] also falls with 2^{32} known-IV chosen texts, due to a piece of blind luck: $K0$ cannot affect $P_0 \oplus C_0$. Similarly, we can obtain attacks requiring 2^{64} known-IV chosen texts and $O(2^{56})$ work against CFB|CFB|CFB^{-1} |CFB^{-1} and OFB[CBC,CBC,CBC]. Our original attack on OFB[CBC,CBC,CBC^{-1}], needing one chosen-IV query and 2^{112} work, can be extended to work with 2^{64} known-IV chosen texts and 2^{112} work. Finally, we can get an attack on CBC|CBC|CBC^{-1} |CBC^{-1} that recovers $K0$ with 2^{64} known-IV chosen texts and 2^{56} trial encryptions, and from there can break the whole quad mode with another $6 \cdot 2^{56}$ trial encryptions and 2^{32} known-IV chosen texts. (These estimates are rough, and the details of the analysis are unchecked.)

Incidentally, these results would disprove Biham's conjectures for several of his modes, *if* we make the major concession of accepting the validity of known-IV attacks. (Additional mild concessions are required in some cases.) For instance, it would show that CBC|CBC|CBC^{-1} |CBC^{-1}, CFB|CFB|CFB^{-1} |CFB^{-1}, and OFB[CBC,CBC,CBC] are not more secure than all triple modes, if we also assume that there is some triple mode which resists all attacks of complexities less than 2^{64}; the latter assumption is quite reasonable, as triple-ECB-DES with outer chaining is an excellent candidate for one such mode. OFB[CFB,CFB,CFB] fall with complexity 2^{56}, which would disprove (if we accept known-IV attacks) Biham's conjecture that it has a security factor of at least 2^{128}. The triple modes (except OFB[CBC,CBC^{-1}]) fall to attacks with 2^{56} complexity (if we accept known-IV attacks), which is less than the conjectured 2^{112} security factor.

These results do not actually refute Biham's conjectured security factors. However, the existence of known-IV attacks of lower-than-expected complexity brings us a step closer to understanding the true security level of these modes.

7 Conclusions

This paper has presented new attacks on all but one of Biham's proposed modes. These attacks rely on the ability to control the IVs, and therefore require quite powerful adversaries which may or may not be a concern in practice. Of all the proposed modes, OFB[CBC,CBC^{-1}] seems to have the best resistance to the chosen-IV attacks which we know of.

These results illustrate the difficulty of building secure modes that contain inner chaining. The danger of internal feedback mechanisms is that the cryptanalyst may be able to probe the internals of the multiple mode of operation by using chosen-text queries; in many cases, this allows the cryptanalyst to isolate the effect of part of the keying material.

This work describes a new failure mode for such systems when the adversary can gain control of IV values. This presents additional evidence for the fragility of constructions based on internal feedback.

We believe that it would be prudent for conservative cryptographic engineers to avoid multiple modes with inner chaining until they are better-understood by researchers. For now, triple-DES-ECB seems to provide more robust—or, at least, better-understood—security.

8 Acknowledgements

The descriptive term "narrow pipe" is due to John Kelsey. The author is deeply grateful to Eli Biham for his comments, which have greatly improved the quality of this work.

References

[Bih94a] E. Biham, "On Modes of Operation," *Fast Software Encryption '93*, LNCS 809, Springer-Verlag, 1994.

[Bih94b] E. Biham, "Cryptanalysis of Multiple Modes of Operation," *ASIACRYPT'94*, LNCS 917, Springer-Verlag, 1994.

[Bih96] E. Biham, "Cryptanalysis of Triple-Modes of Operation," Technion technical report CS 885, 1996.

[CJM97] D. Coppersmith, D.B. Johnson, and S.M. Matyas, "Triple DES Cipher Block Chaining with Output Feedback Masking," *IBM Journal of Research and Development*, vol 40, no 2, 1996.

[OW96] P.C. van Oorschot and M.J. Wiener, "Improving implementable meet-in-the-middle attacks by orders of magnitude," *CRYPTO'96*, pages 228-236, Springer-Verlag, 1996.