CSCI3381-Cryptography

Project Topic: Linear Cryptanalysis of Baby Block Cipher

October 8, 2014

The attack on a block cipher described in another project really worked against a flawed protocol, not against any weakness in the underlying cipher itself. It would work just as well in an attack against an 'ideal' block cipher—a randomly chosen permutation of the set of blocks.

The attack in this project tries to exploit underlying weaknesses in the cipher. Two general methods, called linear and differential cryptanalysis, have been devised for this purpose. For guidance, you should consult 'A Tutorial on Linear and Differential Cryptanalysis', by Howard Heys, at

`http://www.engr.mun.ca/~howard/PAPERS/ldc_tutorial.pdf`

You will experiment with applying the linear cryptanalysis he describes to the Baby Block Cipher.

This cipher was constructed by hunting for permutations of $\{0, \ldots, 15\}$ so that the S-boxes and mixing permutation satisfied certain criteria. In some respects, the result really does resemble a randomly-selected permutation of $\{0,1\}^{16}$ (see the results of the Collision Statistics problem in Assignment 3), but it is likely that any block cipher constructed on the fly this way will deviate from the ideal in other ways that we might be able to exploit.

The idea in linear cryptanalysis is to first look for statistical biases in the S-boxes. For instance, let $x_1, \ldots, x_4$ be the inputs to an S-box and $y_1, \ldots, y_4$ be the corresponding outputs. Pick a subset of the x_i and of the y_i, and XOR them together, for instance

$$x_1 \oplus x_3 \oplus y_2 \oplus y_3 \oplus y_4.$$

We would expect that if the S-boxes were behaving randomly, this expression would have the value 1 for exactly half of the 16 possible input sequences (x_1, x_2, x_3, x_4) and 0 for the other half. In reality, however, there may be significant deviations

from one-half for some choices of the x_i and y_i. Linear cryptanalysis begins by identifying such biases, bootstrapping them into biases in the overall behavior of the block cipher, and then using this information, together with a large amount of known plaintext, to determine bits of the key. The details are given in the Heys paper, which works through a careful example with a cipher very similar to ours. I have not worked through this project myself, but I have gone as far as compiling the biases for the individual S-boxes, and some of them really are quite large. Your job is to see if you can use this to find any bits of the key. This is really a proof-of-concept attack: Since our key is so small, we can easily launch a brute-force attack requiring less effort and less known plaintext than linear cryptanalysis requires. The goal here is to see if the amount of exhaustive search can be reduced significantly by this kind of analysis. (For block ciphers of more realistic size, the effort required for exhaustive search grows exponentially with the size of the blocks and keys, while the cryptanalysis effort grows at a much more modest rate.)

You should begin with reduced-rounds versions of the attack before building up to the full 5 rounds.