Chapter 10

ATTACKS ON MASKING

The use of masking schemes to counteract power analysis attacks is popular for several reasons. For instance, masking can be implemented in software on processors without altering their power consumption characteristics. Probably because of their popularity, many researchers have studied the security of masking schemes and their implementations. It has turned out that virtually every masking scheme can be attacked.

In this chapter, we discuss different types of power analysis attacks on masking schemes, including second-order DPA attacks and template-based DPA attacks. We start by discussing DPA attacks on masking schemes in general. Next, we discuss implementation issues of masking schemes that can be exploited by DPA attacks. Subsequently, we focus on second-order DPA attacks on software implementations. In addition, we explain second-order DPA attacks using templates and template-based DPA attacks on software implementations. Last, we discuss second-order DPA attacks on hardware implementations.

10.1 General Description

Masking provides security against DPA attacks if each masked intermediate value v_m is pairwise independent of the unmasked intermediate value v and the mask m. Hence, only if this pairwise independence does not hold for some reason, a masking scheme is vulnerable to DPA attacks. The DPA attacks, which we have discussed so far, have the property that one intermediate value is predicted and used in the attack. Because only one intermediate value is used, these DPA attacks are also referred to as *first-order* DPA attacks. If several intermediate values are used to formulate the hypotheses, then the corresponding DPA attacks are called *higher-order* DPA attacks. We continue to write DPA attacks in order to refer to first-order DPA attacks in the remainder of this book.

Higher-order DPA attacks exploit the joint leakage of several intermediate values that occur inside the cryptographic device. Remember that due to performance reasons, typical implementations of masking schemes conceal several intermediate values by the same mask. However, even if several masks are used throughout the algorithm, they are generated before the algorithm starts, they are applied to the data and (or) the key, and they are altered by the operations of the algorithm. Consequently, in an implementation where efficiency (memory, speed) is needed, it is always the case that a mask (or a combination of masks) and an intermediate value that is concealed by this mask (or a combination of masks) occur in the device. Hence, in practice it is typically not necessary to study higher-order DPA attacks in general. In practice, it is sufficient to concentrate on higher-order DPA attacks that exploit the leakage that is related to *two* intermediate values. These attacks are called *second-order* DPA attacks. The two intermediate values can either be two values that are concealed by the same mask or a masked value and the corresponding mask.

Second-order DPA attacks exploit the joint leakage of two intermediate values that are processed by the cryptographic device.

10.1.1 Second-Order DPA Attacks

Second-order DPA attacks exploit the leakage of two intermediate values that are related to the same mask. In general, this leakage cannot be exploited directly because the two intermediate values often occur in different operations of the algorithm. Hence, they might be computed subsequently and contribute to the power consumption at different times. In this case, it is necessary to preprocess the power traces in order to obtain power consumption values that depend on both intermediate values.

However, even if the intermediate values contribute to the power consumption at the same time, it is possible that the distribution of the power consumption has the same mean but different variances for all hypotheses. In this case, a DPA attack using the statistical methods that we have used in Chapters 4 and 6 do not succeed because these methods work with the mean value. In order to mount successful DPA attacks in this case it is necessary to either use other statistical methods that exploit the variance, or to preprocess the traces in such a way that the mean-based methods work. The preprocessing is typically done in step 2 of a DPA attack, which consists of recording the power consumption of the device.

Second-order DPA attacks work in the same way as first-order DPA attacks except that they sometimes require preprocessing the power traces.