Application note - rev 1.0 - April 22, 2021

How clock randomization strengthens an SoC's defenses against the threat of side-channel attacks

How to break 'unbreakable' encryption algorithms: the purpose of the side-channel attack

In modern cryptography, the algorithms for encrypting sensitive data such as secret keys are in the public domain; their operation is standardized and well understood. The protection which a secure device enjoys is derived from the data which are unique to each device: the secret keys from which the cryptographic algorithms operate.

Attackers have various reasons for wanting to discover a device's secret keys. State actors such as espionage agencies might want to gain access to sensitive data via a device which provides access to a secure military or government network.

Criminal organizations can deploy considerable resources – both high-tech equipment and high-level expertise – in pursuit of financial gain. Attacks by criminals can include attempts to obtain secret keys so that they can reverse-engineer IP to clone valuable products and sell counterfeit versions on the open market. In the case of a subscription service such as pay TV, the criminal's goal is to sell an illegal device which enables consumers to bypass the service provider's security and payment mechanisms and thus obtain the service for free.

Organized crime also seeks to steal money directly, for instance by compromising a device which provides access to a bank's secure network, or by exposing the secret keys which protect a payment terminal, allowing goods to be 'purchased' while blocking the transfer of money from the criminal buyer to the seller.

In practice, it is almost impossible to steal private keys from SoCs or from their manufacturer. From the SoC through to the end product manufacturer, all participants in the supply chain deploy elaborate physical and logical measures to maintain the integrity of secure production and provisioning processes. Private keys are encrypted using practically unbreakable cryptographic algorithms, and are stored in secure memory space. This means that a properly designed SoC's secrets cannot in practice be directly exposed.

So attackers use a variety of techniques for discovering secrets by analysis or monitoring of the device itself. The side-channel attack (SCA) is one of the most widely used of these techniques.

Why the side-channel attack is an effective technique

The side-channel attack is so called because it seeks to gain information about a data stream, not by monitoring that data stream (a string of 1s and 0s) directly, but by analyzing secondary phenomena which can be correlated to the primary data stream. These secondary phenomena include electro-magnetic emissions, and changes in the power consumed by the processor which executes algorithmic instructions. Side-channel attacks most often use differential power analysis (DPA) of the device under attack.

The sequence of operations of the familiar cryptographic algorithms, such as ECDSA and AES, is well known. These operations have a characteristic power signature which is dependent on the data processed by the algorithm. So by measuring and recording the power signals in a target device, an attacker can discover its secret key. The laboratory equipment required to detect a target's power signature can be readily obtained at a cost which is easily affordable by organized criminal gangs and others.

To perform DPA, attackers repeatedly – many thousands of times – force a sample device through a secured operation such as secure boot, each time recording the current and voltage waveforms. They aggregate these thousands of data samples: by applying sophisticated statistical analyses to them, they can detect correlations which can indicate the timing of cryptographic operations, and potentially reveal the value of the bits which the operation is processing.

The critical importance of timing

An essential feature of a successful side-channel attack using DPA is that the samples have to be precisely and accurately synchronized: the correlations between multiple data samples will only be valid if the start point for each waveform is at exactly the same place in the process under analysis.

In a conventional electronic circuit, synchronization is straightforward, because electronics systems by design operate from a regular timing signal provided either by an integrated oscillator, or through an external device such as a crystal oscillator.

In mounting a side-channel attack which implements DPA, attackers hope that the timing of one instance of a cryptographic operation will be identical to the timing of every other instance, because its timing will be governed by a conventional, regular clock signal.

Building an SoC's defenses against side-channel attack: clock randomization

Embedded device manufacturers cannot prevent attackers from obtaining their product and attempting to subject it to differential power analysis. But the manufacturer of the device's SoC can erect obstacles which make the implementation of an effective DPA attack more time-consuming, expensive and difficult than it would otherwise be. By increasing the effort and investment required to mount an attack, and thus reducing the profit or benefit which might be gained if the attack is successful, these obstacles can help to reduce a device's exposure to the risk of attack. A highly effective way to mount such an obstacle against DPA attack is clock randomization.

Since the synchronization of multiple power analyzer waveforms depends on a clean, regular clock signal, DPA can be made more difficult to implement by deliberately introducing random jitter into the clock signal which governs cryptographic operations. Clock variations which may be applied to disrupt DPA can include:

- Randomly skipping clock cycles, altering the signal's average frequency
- Adding jitter to the clock edges, which randomly modulates the signal's period from edge to edge without modifying its average frequency

When these random types of variation are applied, the shape of each sampled waveform will be uniquely and non-repeatedly distorted in the time domain. This distortion renders the discovery of correlations in the current/voltage patterns between one sample and another much more difficult: attackers are obliged to perform complex, time-consuming re-synchronization before they can even begin the actual DPA itself. For the SoC manufacturer, on the other hand, the implementation of clock randomization is relatively easy – much easier than the development of algorithms intended to harden a system against side-channel attack.

When clock randomization is implemented properly, the scale of the random effects induced in the clock signal is limited to a range which avoids impairing the operation of the SoC's cryptographic algorithms.

At the same time, the impact of the silicon IP which performs clock randomization in an SoC can be very small, both in terms of die area, and in its effect on system power consumption and execution speed. This impact can moreover be minimized by applying clock randomization only to sensitive blocks that perform secure operations. It is normally a simple matter to resynchronize the outputs from secure functional blocks with the rest of the SoC.

INVIA's Secure Clock Generator: clock randomization for secure SoCs

SoC manufacturers can help to secure their devices against the DPA method of implementing side-channel attacks by building clock randomization capability into the silicon. A simple, proven and secure way to do this is with an Invia IP product, the fully integrated Secure Clock Generator.

The Secure Clock Generator provides an on-the-fly, switchable jittered clock with a programmable frequency range. It also produces a secondary clock signal generated by on-the-fly period masking.

As well as performing clock randomization, the Invia IP incorporates a frequency monitor to protect against fault injection attacks aimed at disturbing secure or cryptographic operations by destabilizing the proper clock signal. Fault attacks may be induced by an external disturbance, such as a high-power laser pulse shot into a specific location on the SoC die. To be successful, a laser pulse attack must be highly accurate, both in its location and its timing. The pulse is normally synchronized with signals from the chip under attack. The clock jitter produced by the Invia Secure Clock Generator can make the synchronization required for the timing of the laser pulse much more difficult to achieve.

The Invia IP itself is highly reliable, implementing robust power-up, power-down and standby sequences to avoid the risk of clock glitches. Embedding the Secure Clock Generator imposes low power and area costs. Operating current is less than 120 µA at 120 MHz, and silicon area smaller than 0.05 mm^2 in a typical 55 nm CMOS process.

The IP is silicon-proven in 130 nm, 65 nm and 55 nm CMOS processes, and has been deployed in multiple ICs which are in volume production.