

Research Statement

Antonio Faonio

✌ Computer Science Department,
Aabogade 34, Aarhus N, 8200, Denmark

✉ afaonio@gmail.com,

➡ <https://00af.keybase.pub/>,

☎ { +45 9187 0127, +39 348 7924560 }

My research area is cryptography. From ancient Greek, “crypt” means hidden while “graphy” refers to the process of writing. Once considered the art of “secret communication”, nowadays, cryptography is much more. Rapid technological advances and ubiquitous social and communication networks (such as the Internet) have led to the development of new complex digital services which security requirements go far beyond the need of “communicating secretly”.

In modern terms, cryptography is concerned with the feasibility or infeasibility of securely realizing a task. The answer to whether or not a task can be securely realized depends on the assumed power of the adversary and the network model. The goal is to obtain security while minimizing the need for trust. The obtained security holds in a strong mathematical sense, meaning that breaking a secure cryptographic protocol is either impossible or it would require to solve some computational problem which is believed to be hard with the current technology.

In what follow I describe the main themes concerning my research, together with directions for future research.

Research

The achievements of my research are mostly on leakage resilient cryptography, but not limited to this. In fact, I participated to research projects in different topics of cryptography, both from applied to theoretic prospective.

Leakage and Tamper Resilient Cryptography

The areas of research of leakage resilience can be motivated with a very simple observation, which has made and makes me very enthusiastic to work on these fields: A cryptographic primitive is eventually implemented either as a program or as a circuit and will be surrounded by a specific environment. An attacker can take advantage of implementation-specific characteristics, and indeed, provably secure systems turned out to be physically vulnerable. This vulnerabilities are exploited both via side-channel and tampering attacks.

Standard security definitions for cryptographic primitives (e.g., encryption or signature schemes) assume implicitly that the adversary has only black-box access to the underlying algorithms, however side-channel attacks have shown that this is not the case, in fact, auxiliary information about the secret state can be gathered, for example, through analysis of the magnetic field, by probing wires and even via acoustic noise.

The scenario is even more complicated in the case of tampering attacks. These are *active* attacks, which generally rely on inducing fault on the real-world implementations (for example by microwaving the device or exposing to heat or EM radiation) and allow in some case to completely break otherwise provably secure schemes.

In the recent years a lot of progress has been done, both in term of practical countermeasures on specific side-channel and tampering attacks and in term of realizing cryptographic primitives provably secure in gradually stronger models.

My research lies in the second trend, with the goal of making efficient cryptographic schemes secure in reasonable models which capture existing side-channels and tampering attacks.

Here, I would like to stress on the words “reasonable” and “efficient”. The common critique is whether the known theoretical models cover physical leakages in an appropriate way or not. As example, while the bounded leakage model (probably the most known and studied leakage model) offers a beautiful abstraction by only bounding the nature of the leakage on its output size, a power measurement attack can hardly be described by a few bits of information. The second critique is that leakage-resilient tamper-resilient schemes are still far from being practical. For the reasons described above, my works concentrate on general models with a stress on efficiency. With my co-authors I contributed to the development of leakage-resilient primitives.

- In [13], we developed signature schemes in the noisy and fully leakage model. Here we considered a leakage model where the only constraint is that the secret material (both the secret key and the randomness of the signing process) must retain some uncertainty (in an information-theoretic sense) after the leakage. By considering attacks also on the randomness we push the boundary even further, in fact, it is natural that an enemy capable of side-channel attacks would also leak from the random source. The standard notion of signature scheme does not cope well with leakage, short signature schemes are indeed impossible according to the standard definition, however, we were able to provide secure schemes with short signature size (and indeed better efficiency, at least in term of communication complexity) by relaxing the definition to “the best we can hope for”.

Namely, we relaxed the definition considering an adversary that might be able to produce some signatures (this because, in principle, it could leak them), however it cannot produce more signatures that it can leak. The paper also contains a fairly practical scheme using the random oracle model, the scheme achieves the extra efficiency of (a noisy version of) the bounded retrieval model (BRM) where the secret key can be as big as we want without any degradation of all the other efficiency parameters.

I consider the BRM very interesting: having a 1GB long secret key stored in our device, right now, is not a big deal, on the other hand, performing a side-channel attack that stealthily leaks all this amount of data (either using side-channel or a by infecting the device) seems to be quite unpractical. The bounded retrieval model is the subject of another paper that I co-authored:

- In [5], we showed that Proof of Storage (specifically, their zero-knowledge version) can be proved to be secure identification scheme in the BRM. This led to a very practical RSA based scheme, which I could easily implement on top of an already implemented RSA Proof of Storage [4].

I got interested in general purpose compilers that provide leakage resilience to any cryptographic functionality.

- In [12], we analyzed the tool of Leakage Resilient Codes (LRC) that are very useful in the setting I mentioned before (see Dziembowski and Faust [9]). One of the drawbacks of using LRC is that the inputs need to be encoded in a leak-free environment. We investigated if this is necessarily the case. Unfortunately, we showed that LRC that are also resilient to small leakage from the encoding process (which we dubbed Fully Leakage Resilient Codes) are impossible to achieve. This is the case because the leakage function can be very complicated. Fortunately, we showed that the impossibility result can be bypassed in certain circumstances e.g., by putting some limitations on the complexity of the leakage or by having milder setup assumptions like the common reference string.

More recently, I have been involved in a project which deals with tamper resilience.

- In [15], we showed efficient constructions of secure signature scheme and public key encryption scheme against adversaries that can arbitrarily tamper but only for a limited amount of time with the memory of the device. Notice that there are strong negative result in the case of an unbound number of tampering attempts, on the other hand, the bounded tampering model seems very reasonable in practice. Intuitively, the more the attacker tampers with the memory the more likely the device will break down and stop to work. A previous result of Damgaard et al. [8] showed a construction of a signature scheme in the random oracle model and public key encryption relying on expensive pairing operations. We obtained the first construction of signature scheme in the standard model and the first public-key encryption construction free of pairing (avoiding non-interactive zero-knowledge proofs). Our constructions are based on generic building blocks, and admit efficient instantiations under fairly standard number-theoretic assumptions.

From my point of view, it is necessary to find the right balance between efficiency and a good approximation of the real capability of the adversary: standard model primitives are very fast but they cannot be considered truly secure. Bounded leakage is a halfway through, it is not quite the answer against real side-channel attacks, but it is a significant step forward and fairly efficient schemes exist. On the other hand, the landscape for more general models e.g., the noisy leakage model or the fully leakage model is still broadly unknown. Finally, while tampering with the memory of the device has been extensively studied, and we showed that efficient solutions exist in the bounded tampering model, less it has been done for more general forms which involve, for example, tampering with the computational model.

Other works on Theoretic and Applied Cryptography

Cryptography is an awesome field to work on, in fact, while motivated by highly practical problems, it allows to originate new surprising mathematical concepts that were never studied before. Below I give some examples from my research achievements in foundational cryptography.

- In [3], we started with the practical question of making the BlockChain¹ technology environmental friendly². The paper introduces the concept of Proof of Space (PoSpace). Briefly, a PoSpace protocol allows to succinctly show to a verifier the posses of some space, specifically, to compute such a proof, the prover must use a specified amount of space. To compare, we proposed an alternative to the Proof-of-Work technology which would make the Bitcoin system secure, by using space (like unused hard drive space) rather than computational power. Apart of proposing a new concept, we showed constructions in the Random Oracle Model, both interactive and non-interactive.
- In [14], we initiated the characterization of proof systems that we called Predictable Arguments of Knowledge (PAoK). The project started when we were working on the paper of Fully Leakage Resilient Codes, we noticed that, in order to strengthen our result, we needed a new form of proof system. Soon we realized that the concept deserved more interest. Specifically, we analyzed argument of knowledge³ for NP where the verifier can predict the answer that the prover will give. Thanks to this powerful tool we were able to show that any LRC (in the popular split-state model) can be broken by leaking only 7 bits from the encoding process. But this is only an application that shows the power of PAoK. More interestingly, we showed connections with other cryptographic primitives, we showed that PAoK can be made extremely succinct (one round and one bit from the prover to the verifier are enough!), we gave a full characterization of its zero-knowledge versions and analyzed its power in the context of adaptive security. This is, up to now, probably the most theoretic work which I was involved in.

¹The BlockChain is the public ledger used by Bitcoin, which is a peer-to-peer electronic cash system.

²See “Virtual Bitcoin Mining Is a Real-World Environmental Disaster” by Mark Gimein.

As mentioned already, I am open to do research in all the aspects of cryptography. I am also interested in applied cryptography:

- In [6], we described an extension of the Bitcoin protocol that preserves its decentralized nature, while also enabling payers to optionally specify the involvement of a trusted authority that attests to the identity of the payee, by requiring payees to use certified Bitcoin addresses. More specifically, we introduce the concept of Bitcoin addresses that need to be generated with the support of a semi-trusted authority. Notice that, while the semi-trusted authority can mint coins on behalf of a particular user, it cannot spend any of them. Nicely, these certified addresses are allowed to co-exist with the standard auto-generated Bitcoin addresses and can be easily integrated without any change in the specifications of the BlockChain.

Future Directions

I expect to continue working in the areas I mentioned in the future, while remaining open to new opportunities. Below are just a few examples to illustrate the kind of problems I expect to pursue in the coming years.

The concept of Proof of Space has received a lot of interest (e.g., see [1, 2, 10, 16] for related works³), however, all the proposed constructions of Proof of Space or of related cryptographic primitives are in the Random Oracle Model. It is still an open problem to find a protocol in the standard model. I have in mind to leverage some recent advance in complexity (see Raz [17]) to tackle the problem.

Currently, I am involved in research projects that spread from tamper resilient cryptography to subversion resilience. An interesting research trend in tamper resilience is the fascinating tool of Non-Malleable Code (NMC). Since their introduction by Dziembowski et al. [11], NMC received a lot of interest. Devise new flavor of NMC always brings an improvement in the context of tamper resilience. I have in mind several extensions, for example, recently, I'm working on non-malleable and leakage resilient codes in the split-state model that can be succinctly updated. In this model the codeword is divided in two (or more) pieces and the tampering happens independently in each piece. We found out that in the information theoretic setting, the updating procedure necessarily needs interaction between the pieces. However, in computational setting, this does not seem to be the case. In fact, we have already some feasibility results that need to be further investigated.

In the revelations of Edward Snowden have evidenced that some cryptographic protocol specifications, for example the standard Dual EC in NIST SP 800-90A, were modified to embed backdoors. Recently, the ability of substituting a cryptographic algorithm with an altered version was (re)considered by Bellare et al. [7] (see also [18]). The landscape shown is mostly negative: all stateless and randomized encryption schemes can be stealthily subverted by a powerful adversary, a big brother, in such a way that no privacy can be assured. In [7] is shown a class of encryption schemes that are subversion resilient, however, we have a preliminary result showing that in a more general setting where the subversion-resistant encryption scheme is triggered by an application, that in turn can be subverted, subversion resilience is not enough to assure privacy. Currently I am working on way to overcome this negative result, the idea is to apply knowledge from leakage and tamper resilience cryptography in this context.

In my view, privacy and security are essential for a free and functioning society. Doing research in cryptography is therefore not only real fun but also an active way to collaborate for a better world. Leakage and tampering attacks are getting cheaper and easier to implement while substitution attacks were showed to be a real threat. In my opinion, it is, therefore, important to bring to a practical level both leakage and tamper resilient schemes, and actively get involved with the research on the new and still undiscovered setting of subversion resilience.

³In particular, [10] is an independent work.

References

- [1] Joël Alwen, Binyi Chen, Chethan Kamath, Vladimir Kolmogorov, Krzysztof Pietrzak, and Stefano Tessaro. On the complexity of script and proofs of space in the parallel random oracle model. In *EUROCRYPT, Part II*, pages 358–387, 2016.
- [2] Joël Alwen and Vladimir Serbinenko. High parallel complexity graphs and memory-hard functions. In *STOC*, pages 595–603, 2015.
- [3] Giuseppe Ateniese, Ilario Bonacina, Antonio Faonio, and Nicola Galesi. Proofs of space: When space is of the essence. In *SCN*, pages 538–557, 2014.
- [4] Giuseppe Ateniese, Randal C. Burns, Reza Curtmola, Joseph Herring, Lea Kissner, Zachary N. J. Peterson, and Dawn Xiaodong Song. Provable data possession at untrusted stores. In *CCS*, pages 598–609, 2007.
- [5] Giuseppe Ateniese, Antonio Faonio, and Seny Kamara. Leakage-resilient identification schemes from zero-knowledge proofs of storage. In *IMACC*, pages 311–328, 2015.
- [6] Giuseppe Ateniese, Antonio Faonio, Bernardo Magri, and Breno de Medeiros. Certified bitcoins. In *ACNS*, pages 80–96, 2014.
- [7] Mihir Bellare, Kenneth G. Paterson, and Phillip Rogaway. Security of symmetric encryption against mass surveillance. In *CRYPTO, Part I*, pages 1–19, 2014.
- [8] Ivan Damgård, Sebastian Faust, Pratyay Mukherjee, and Daniele Venturi. Bounded tamper resilience: How to go beyond the algebraic barrier. In *ASIACRYPT, Part II*, pages 140–160, 2013.
- [9] Stefan Dziembowski and Sebastian Faust. Leakage-resilient circuits without computational assumptions. In *TCC*, pages 230–247, 2012.
- [10] Stefan Dziembowski, Sebastian Faust, Vladimir Kolmogorov, and Krzysztof Pietrzak. Proofs of space. In *CRYPTO, Part II*, pages 585–605, 2015.
- [11] Stefan Dziembowski, Krzysztof Pietrzak, and Daniel Wichs. Non-malleable codes. In *ICS*, pages 434–452, 2010.
- [12] Antonio Faonio and Jesper Buus Nielsen. Fully leakage-resilient codes. *IACR Cryptology ePrint Archive*, 2015:1151, 2015.
- [13] Antonio Faonio, Jesper Buus Nielsen, and Daniele Venturi. Mind your coins: Fully leakage-resilient signatures with graceful degradation. In *ICALP, Part I*, pages 456–468, 2015.
- [14] Antonio Faonio, Jesper Buus Nielsen, and Daniele Venturi. Predictable arguments of knowledge. *IACR Cryptology ePrint Archive*, 2015:740, 2015.
- [15] Antonio Faonio and Daniele Venturi. Efficient public-key cryptography with bounded leakage and tamper resilience. *IACR Cryptology ePrint Archive*, 2016:529, 2016.
- [16] Sunoo Park, Krzysztof Pietrzak, Albert Kwon, Joël Alwen, Georg Fuchsbauer, and Peter Gazi. Spacemint: A cryptocurrency based on proofs of space. *IACR Cryptology ePrint Archive*, 2015:528, 2015.
- [17] Ran Raz. Fast learning requires good memory: A time-space lower bound for parity learning. *IACR Cryptology ePrint Archive*, 2016:170, 2016.
- [18] Adam L. Young and Moti Yung. Kleptography: Using cryptography against cryptography. In *EUROCRYPT*, pages 62–74, 1997.