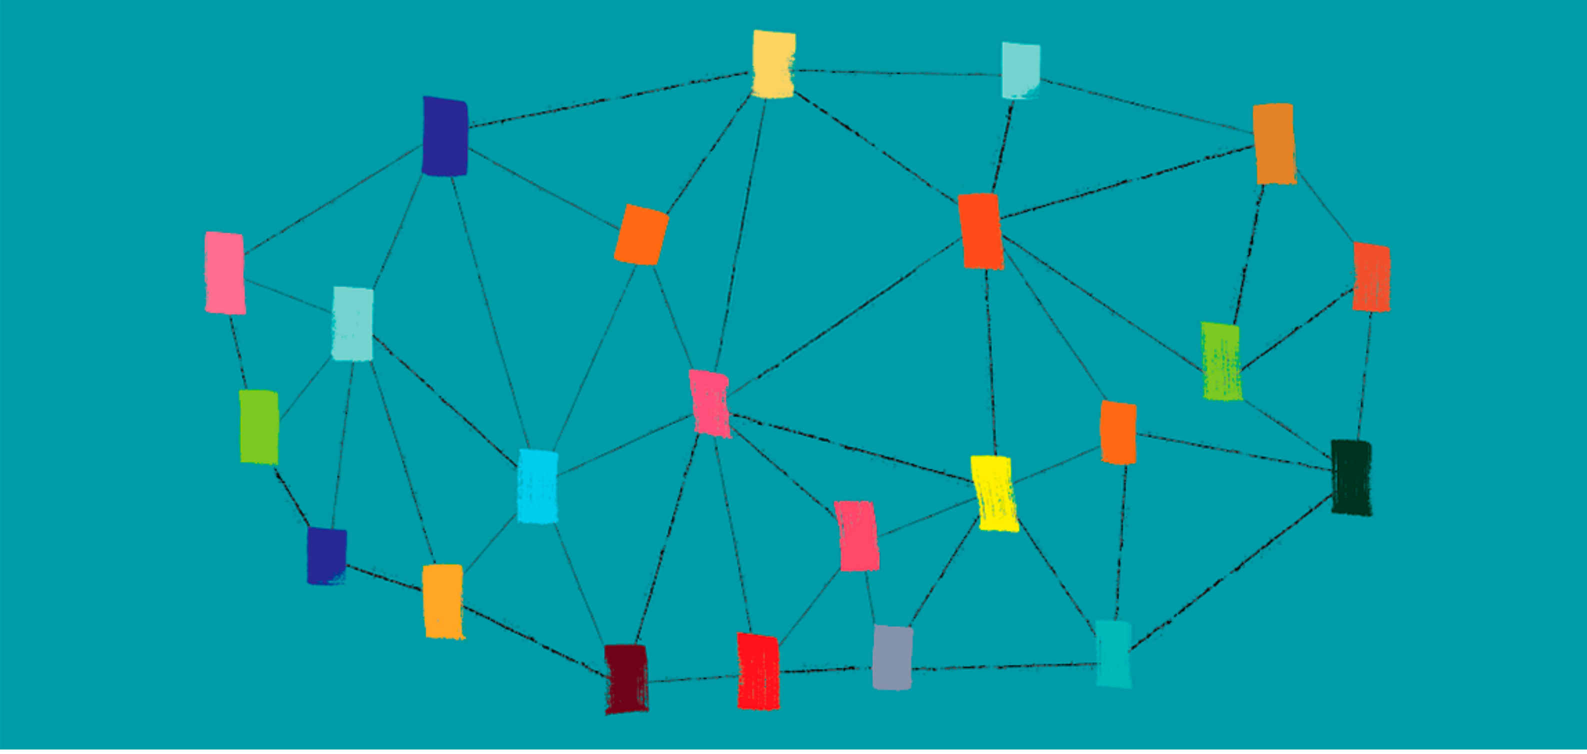

Workshop on
Privacy preserving systems, software and tools

Organized by the
Laboratory of Cryptography and Cybersecurity
Department of Mathematics and Physics
Roma Tre University

October 24, 2022
Aula Urbano VIII, Argiletum - Roma Tre University
Via della Madonna dei Monti 40

Registration required at: http://www.matfis.uniroma3.it/cryptoworkshop/

Chair **Marco Pedicini,** Roma Tre University

9:00 – 9:10 a.m.
Marco Cianfriglia, Roma Tre University
Opening

9:10 – 9:50 a.m.
Massimiliano Sala, University of Trento
On the equivalence of two post-quantum cryptographic families

9:50 – 10:30 a.m.
Christian Mouchet, EPFL
Computing across Trust Boundaries with Multiparty Homomorphic Encryption and the Lattigo library

10:30 – 11:00 a.m.
Coffee break

11:00 - 11:40 a.m.
Daniele Friolo, La Sapienza University of Rome
Multi-Key and Multi-Input Predicate Encryption from Learning With Errors

11:40 – 12:20 a.m.
Dario Pasquini, EPFL
Privacy-Preserving Collaborative Machine Learning ?

12:20 a.m. - 1:00 p.m.
Sinem Sav, EPFL
Privacy-Preserving Federated Learning with Multiparty Homomorphic Encryption

1:00 – 1:10 p.m
Closing

DMF
ROMA TRE
UNIVERSITÀ DEGLI STUDI

ABSTRACTS:

Massimiliano Sala

On the equivalence of two post-quantum cryptographic families

The maximum likelihood decoding problem (MLD) is known to be NP-hard and its complexity is strictly related to the security of some post-quantum cryptosystems, that is, the so-called code-based primitives. Analogously, the multivariate quadratic system problem (MQ) is NP-hard and its complexity is necessary for the security of the so-called multivariate-based primitives. In this paper we present a closed formula for a polynomial-time reduction from any instance of MLD to an instance of MQ, and viceversa. We also show a polynomial-time isomorphism between MQ and MLD, thus demonstrating the direct link between the two post-quantum cryptographic families.

Christian Mouchet

Computing across Trust Boundaries with Multiparty Homomorphic Encryption and the Lattigo library

Abstract:Recent Multiparty Homomorphic Encryption (MHE) techniques have given rise to an highly efficient family of secure multiparty computation (MPC) solutions that are increasingly employed for concrete, real-world applications. One reason for this emergence is that MHE-based MPC solutions address several short comings that are usually associated with secure computation. Notably, they have a low communication complexity and they support delegation to third parties, yet without relying on trust-delegation or non-collusion assumptions. As such, they enable large-scale MPC even among computationally weak or sporadically online clients. Moreover, MHE-based approaches directly benefit from there markable progress in the efficiency and useability (single-party) HE schemes. Hence, MHE techniques are bound to play an increasingly important role in the next generation of secure computing systems. In this talk, I give an over view of the current MHE techniques and their use in MPC. Then, I present our efforts to make them accessible to researchers and practitioners through the Lattigo open-source library. Finally, I cover some of the challenges and future research directions in the MHE-based MPC.

Daniele Friolo

Multi-Key and Multi-Input Predicate Encryption from Learning With Errors

In Predicate Encryption (PE) a user encrypts a message with a predicate input. The resulting ciphertext can be later decrypted only by users holding a key related to a predicate satisfying the input used during encryption. Differently from Attribute-Based Encryption (ABE), the privacy of the predicate input (attribute) is maintained.

In our work, we put forward two natural generalizations of predicate encryption (PE) dubbed multi-key and multi-input PE.

More in details, our contributions are threefold.

- Definitions. We formalize security of multi-key PE and multi-input PE following the standard indistinguishability paradigm, and modeling security both against malicious senders (i.e., corruption of encryption keys) and malicious receivers (i.e., collisions).
- Constructions. We construct multi-key and multi-input PE supporting the conjunction of poly-many arbitrary single-input predicates, assuming the hardness of the standard learning with errors (LWE) problem.
- Applications.

We show that multi-key and multi-input PE for expressive enough predicates suffices for interesting cryptographic applications, including matchmaking encryption (ME) and non-interactive multi-party computation (NI-MPC). As a corollary, plugging in our concrete constructions of multi-key and multi-input PE, we obtain the first construction of ME for arbitrary policies, as well as NI-MPC with partial re-usability for all-or-nothing functions and a constant number of parties, under the standard LWE assumption. Prior to our work, all of these applications required much heavier tools such as indistinguishability obfuscation or compact functional encryption.

Dario Pasquini

Privacy-Preserving Collaborative Machine Learning ?

This talk is about inaccurate assumptions, unrealistic trust models, and flawed methodologies affecting current collaborative machine learning techniques. In the presentation, we cover different security issues concerning both emerging approaches and well-established solutions in privacy-preserving collaborative machine learning. We start by discussing the inherent insecurity of Split Learning and Fully-Decentralized Machine Learning. Then, we talk about the soundness of current Secure Aggregation protocols in Federated Learning, showing that those do not provide any additional level of privacy to users in the malicious setting. Ultimately, the objective of this talk is to highlight the general errors and flawed approaches we all should avoid in devising and implementing "privacy-preserving collaborative machine learning".

Sinem Sav

Privacy-Preserving Federated Learning with Multiparty Homomorphic Encryption

Training accurate and robust machine learning models requires a large amount of data that is usually scattered across data silos. Sharing or centralizing the data of different healthcare institutions is, however, unfeasible or prohibitively difficult due to privacy regulations. We address the problem of privacy-preserving training and evaluation of neural networks in an N-party, federated learning setting. We propose a novel system, POSEIDON, the first of its kind in the regime of privacy-preserving neural network training. It employs multiparty lattice-based cryptography to preserve the confidentiality of the training data, the model, and the evaluation data, under a passive-adversary model and collusions between up to N−1 parties. To efficiently execute the secure back propagation algorithm for training neural networks, we provide a generic packing approach that enables Single Instruction, Multiple Data (SIMD) operations on encrypted data. Our experimental results show that POSEIDON achieves accuracy similar to centralized or decentralized non-private approaches and that its computation and communication overhead scales linearly with the number of parties. We improve POSEIDON and demonstrate its applicability on biomedical analysis for disease-associated cell classification with single-cell analysis. For this, we design a system, PriCell, for training a published state-of-the-art convolutional neural network in a decentralized and privacy-preserving manner. We compare the accuracy achieved by PriCell with the centralized and non-secure solutions and show that PriCell guarantees privacy without reducing the utility of the data.