Practical Multi-Key Homomorphic Encryption for Efficient Secure Federated Average Aggregation

Financiado por la Unión Europea NextGenerationEU
Plan de Recuperación, Transformación y Resiliencia
GOBIERNO DE ESPAÑA
MINISTERIO DE CIENCIA E INNOVACIÓN
FEDER - FONDO EUROPEO DE DESENVOLVEMENTO REXIONAL "Unha maneira de facer Europa"
UNIÓN EUROPEA
TRUMPET
XUNTA DE GALICIA

A. Pedrouzo-Ulloa[1] A. Boudguiga[2] O. Chakraborty[2] R. Sirdey[2] O. Stan[2] M. Zuber[2]
[1] apedrouzo@gts.uvigo.es, [2] name.surname@cea.fr

atlanTTic research center for Telecommunication Technologies
Universida de Vigo
cea list

6th HomomorphicEncryption.org Workshop
23-24 March 2023, Seoul (South Korea)

Optimizing HE for Federated Average Aggregation

Federated Learning: Many works address the problem of secure aggregation in FL [1]. However, to the best of our knowledge, HE has not been yet fully optimized for this setting.

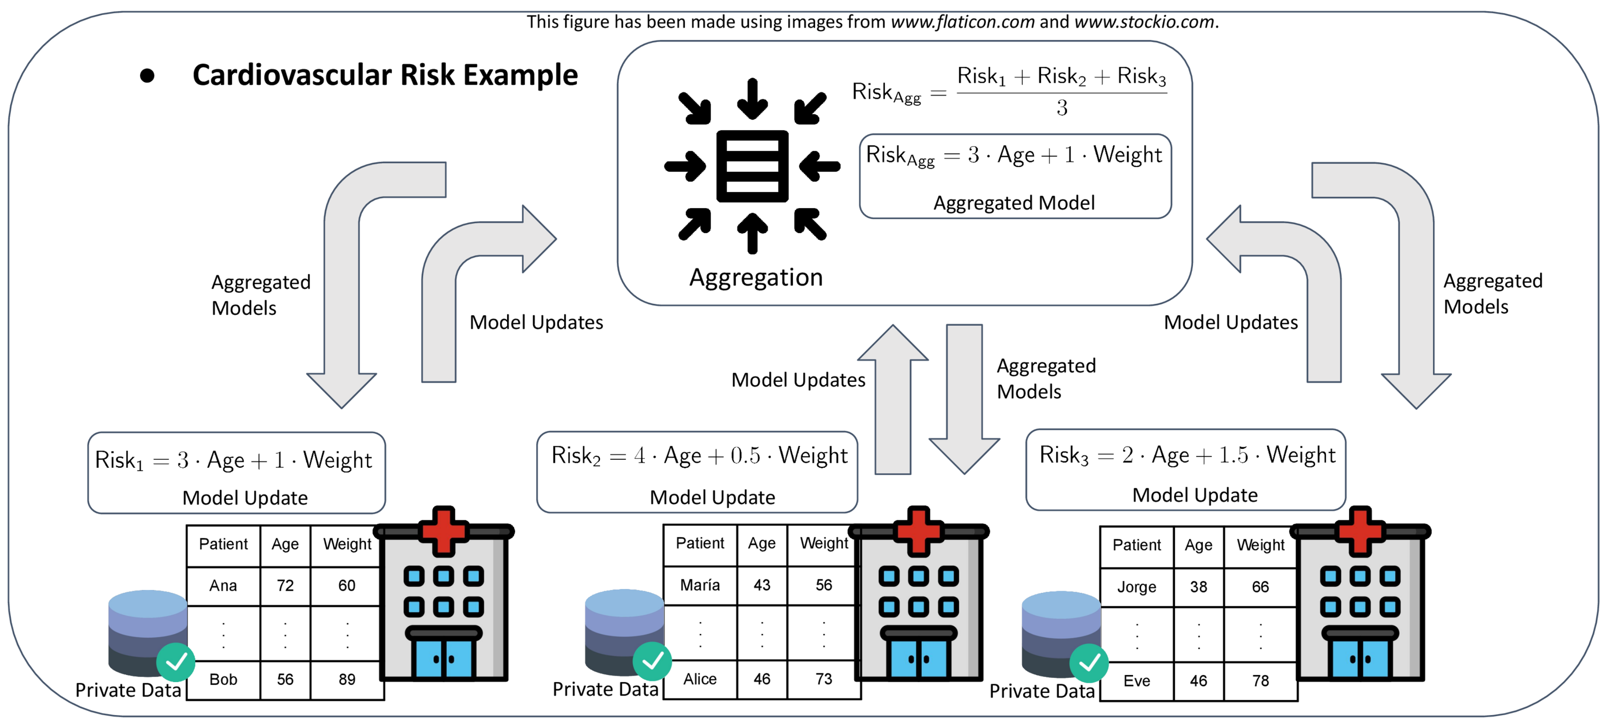

Main objective: Tailor and optimize HE constructions for secure average aggregation.
Main contribution: A lightweight communication-efficient multi-key approach suitable for the Federated Averaging rule [2].

- Communication cost per party is reduced approximately (1) by a half with RLWE, and (2) from quadratic to linear in terms of lattice dimension if considering LWE.
- Secure against malicious aggregators by at most doubling communication cost per party.

Some limitations of current HE-based solutions

Non-Colluding Assumption: Single-Key HE [3] imposes a non-colluding assumption between the aggregator and the owner of the secret key SK.

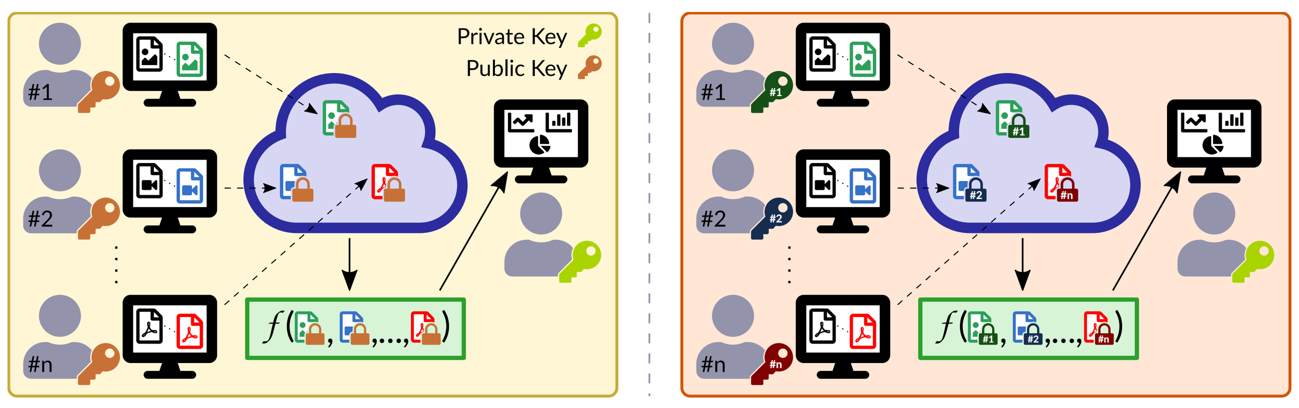

Public Keys: Both Single-Key HE [3] and Threshold HE [4] give access to encryptions of zero (i.e., PK = Enc(0)) under the global secret key SK.

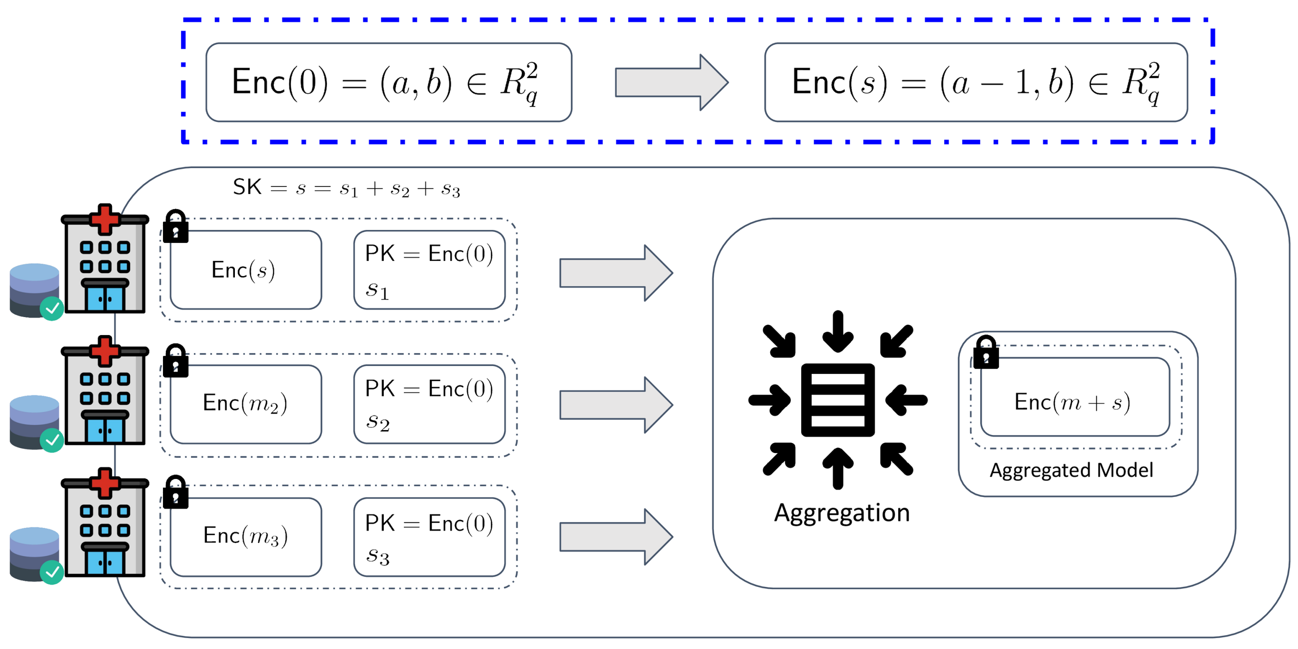

Dishonest Data Owners: A dishonest Data Owner (DO) could easily generate a valid encryption of the global secret key by only having access to the PK.

An upgrade to malicious aggregators

Limiting ciphertexts' malleability: By assuming the Common Reference String (CRS) model, a different "a" term is fixed among all Data Owners during each aggregation round.

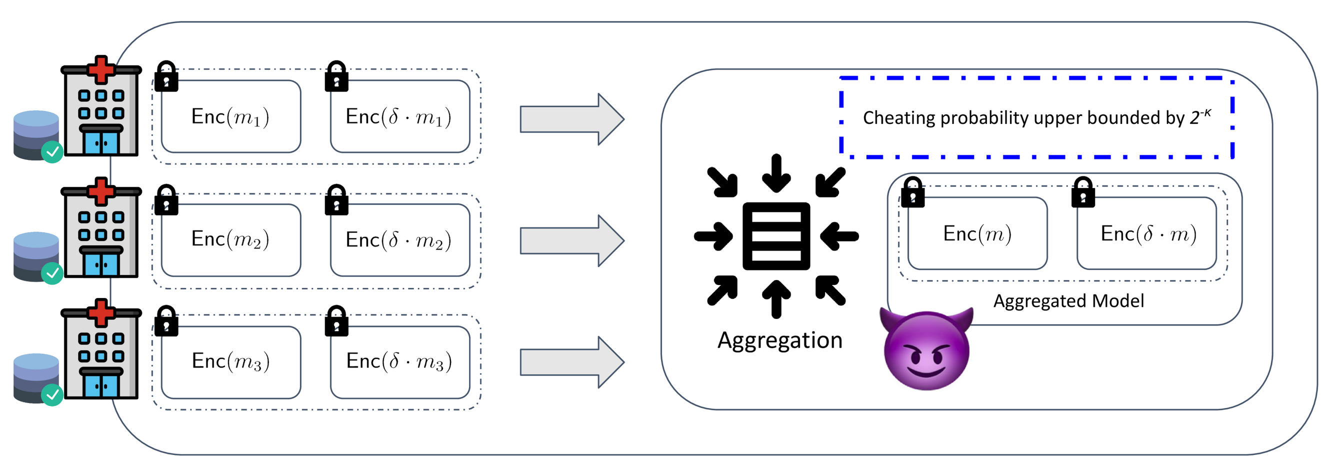

- The Aggregator can only apply additive transformations without being detected.
- An extra condition check can be embedded into Secret-Key ciphertexts (e.g., $\delta \cdot m$ with δ unknown to aggregator). This verifies the honest behavior during aggregation.

Proposed HE-based Protocol

High-level view: Our HE-based protocol for secure aggregation. See [2] for more details.

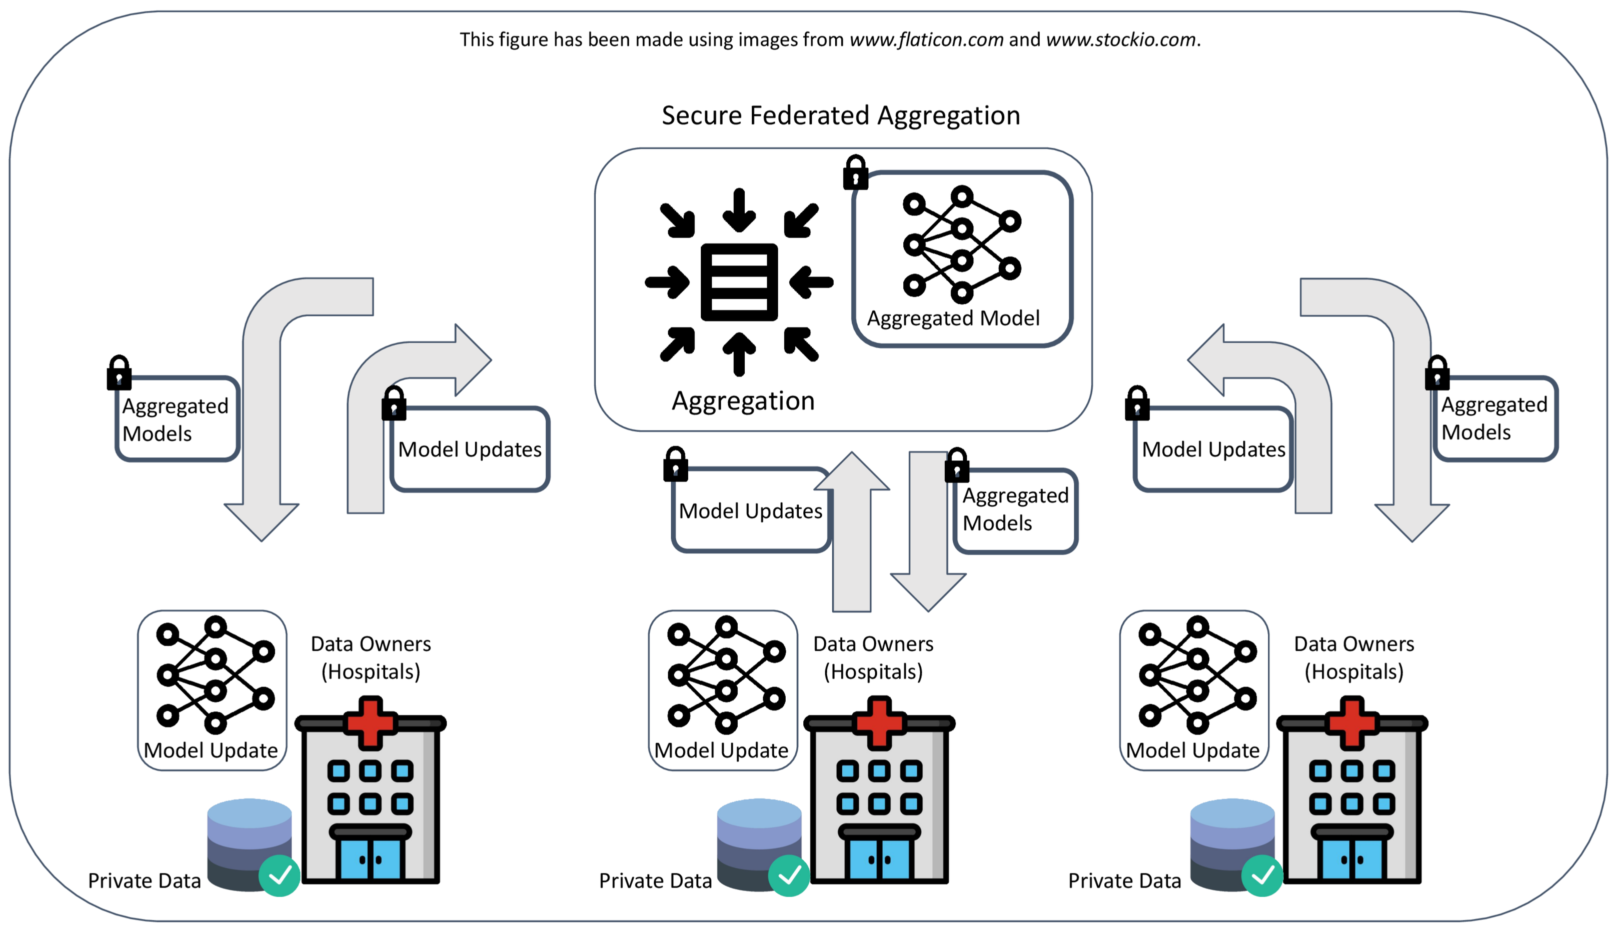

Protocol setup:

- In the CRS model, DOs have access to a common uniformly random a per round.
- All DOs have access to one random polynomial share of zero: $\mathsf{share}_i = r^{(i)}$.

Workflow for a round of our secure aggregation protocol (semi-honest example):

1. DOs encrypt their inputs: The i-th DO ($\forall i$) encrypts its model update m_i as:

$$b_i = a(s_i + r^{(i)}) + e_i + q/p \cdot m_i.$$

2. Aggregation step: $b = \sum_i b_i = a(s + \underbrace{\sum_i r^{(i)}}_{0}) + e = a \underbrace{s}_{\sum_i s_i} + \underbrace{e}_{\sum_i e_i} + q/p \cdot \underbrace{m}_{\sum_i m_i}$.

   Finally, the aggregator sends back $\mathsf{share}^{(\mathsf{agg})} = \lfloor b \rceil_{p'}$ to the DOs.

3. Distributed decryption:

   (a) The i-th DO ($\forall i$) computes $\mathsf{share}^{(i)} = \lfloor a s_i \rceil_{p'}$ and makes it available to all DOs.

   (b) All DOs compute $\left\lfloor \mathsf{share}^{(\mathsf{agg})} - \sum_i \mathsf{share}^{(i)} \right\rceil_p$.

Comparison with other solutions

Next table compares our work with a representative set of HE and MPC solutions.

M: Model Size N: Number of DOs n: lattice dimension M ≈ constant · n	**Ours [2]**	**[5]**	**[3]**	**[4]**	**[6]**
Agg. Comp. Cost	$O(MN)$ add.	$O(MN)$ mult.	$O(MN)$ add.	$O(MN)$ add.	$O(MN^2)$
DO Comp. Cost	LWE: $O(Mn)$ mult. RLWE: $O(M \log M)$ mult.	$O(M)$ exp.	$O(M \log M)$ mult.	$O(M \log M)$ mult.	$O(MN + N^2)$
Total Com. Cost	$O(MN)$	$O(MN)$	$O(MN)$	$O(MN)$	$O(MN + N^2)$
Multiple Keys	✅	🚫	🚫	✅	✅
Passive parties	✅	✅	✅	✅	✅
Malicious Agg.	✅ Verify Agg.	✅ Verify Agg.	🚫	🚫	✅ only DOs input privacy if $T > N/2$
Assumptions	LWE/RLWE	Paillier	RLWE	RLWE	T non-colluding DOs
Flexible Dec.	✅ only DOs contributing to aggregated model	🚫	🚫	🚫	✅ required T out of N DOs

- **HE-based aggregation**: We include RLWE-based Single-Key [3] and Multi-Key [4] schemes. Also Paillier with verifiable computation for malicious aggregators [5].
- **MPC-based aggregation**: We include a work [6] relying on Shamir's Secret Sharing.

References

[1] Mohamad Mansouri, Melek Önen, Wafa Ben Jaballah, and Mauro Conti, "Sok: Secure aggregation based on cryptographic schemes for federated learning," *Proc. Priv. Enhancing Technol.*, vol. 2023, no. 1, pp. 140–157, 2023.

[2] Alberto Pedrouzo-Ulloa, Aymen Boudguiga, Olive Chakraborty, Renaud Sirdey, Oana Stan, and Martin Zuber, "Practical multi-key homomorphic encryption for more flexible and efficient secure federated aggregation (preliminary work)," *IACR Cryptol. ePrint Arch.*, p. 1674, 2022.

[3] Arnaud Grivet Sébert, Renaud Sirdey, Oana Stan, and Cédric Gouy-Pailler, "Protecting data from all parties: Combining FHE and DP in federated learning," *CoRR*, vol. abs/2205.04330, 2022.

[4] Christian Mouchet, Juan Ramón Troncoso-Pastoriza, Jean-Philippe Bossuat, and Jean-Pierre Hubaux, "Multiparty homomorphic encryption from ring-learning-with-errors," *Proc. Priv. Enhancing Technol.*, vol. 2021, no. 4, pp. 291–311, 2021.

[5] Abbass Madi, Oana Stan, Aurélien Mayoue, Arnaud Grivet-Sébert, Cédric Gouy-Pailler, and Renaud Sirdey, "A secure federated learning framework using homomorphic encryption and verifiable computing," 2021, pp. 1–8.

[6] Kallista A. Bonawitz, Vladimir Ivanov, Ben Kreuter, Antonio Marcedone, H. Brendan McMahan, Sarvar Patel, Daniel Ramage, Aaron Segal, and Karn Seth, "Practical secure aggregation for privacy-preserving machine learning," in *ACM SIGSAC CCS*. 2017, pp. 1175–1191, ACM.