PhD Proposal. Towards Differentially Private SQL Query Interpretation: A Comprehensive Approach and Implementation in PostgreSQL.

Jean-François Couchot[*], Catuscia Palamidessi[†], Veronika Sonigo[‡]
couchot@femto-st.fr, catuscia@lix.polytechnique.fr, vsonigo@femto-st.fr

April 8, 2024

1 Introduction

The European Union aims to hold data processors accountable by requiring that critical algorithmic decisions made by AI are explainable, transparent, and safe, as formalized in [Com21]. Achieving explainability relies on the ability to thoroughly evaluate AI, which often involves real database data. Database Management Systems (DBMS) can handle personal data, but should only disclose it to legitimate recipients in line with GDPR regulations. Encryption and pseudonymization are essential, but not sufficient for data protection. Anonymizing data to prevent re-identification is crucial in the era of big data, with approaches including syntactic methods and differential privacy (DP)[DR14], the latter being widely accepted for its probabilistic measurement of information leakage. However, DP's complexity and noise introduction may compromise utility if not carefully managed.

The aim of this PhD is to develop and assess a "privacy preserving" method for interpreting SQL queries, focusing on ensuring differential confidentiality within PostgreSQL. These queries will cover various operations, from basic Select-Project-Join-Aggregation (SPJA) forms to exporting segments of the database (DUMP). This project builds upon the PostgreSQL Anonymizer tool [Clo23] by DALIBO, extending its existing anonymization features to include models compliant with differential privacy (DP), along with some modifications.

2 Working plan

It is divided into three steps.

2.1 From state of the art DP-based query SQL query sanitization to their implementation in PostgreSQL Anonymizer

The goal of this subtask is to study and then to implement in the PostgreSQL Anonymizer the most promising (in terms of utility) approaches of query sanitization [JNS18, KTH$^+$19, WZL$^+$20, TGMR22, TGMR22, DFY$^+$23, GHIM19]. Notice that many codes of these theoretical works are available online allowing a "rapid" implementation.

A Summary of Knowledge is planned to be submitted at 6 months, and the implementation will be made available at 12 months.

[*] https://members.femto-st.fr/jf-couchot/en/research

[†] http://www.lix.polytechnique.fr/~catuscia/

[‡] https://members.femto-st.fr/veronika-sonigo/

2.2 Reduction of composing mechanisms in DP relaxations

While maintaining high privacy protection for individual queries in a database is straightforward, the real-world scenario involves executing multiple SPJA queries consecutively. Therefore, the focus shifts to addressing the composition of query-answering mechanisms. Existing techniques like the Gaussian mechanism relaxation offer some resistance to composition effects but are not entirely immune. Recently, other relaxations of differential privacy, such as Renyi-DP[Mir17] or concentrated DP [DR16], have emerged, particularly suited for handling compositions, often found in iterative Machine Learning approaches.

The primary goal of this step is to investigate how the interpretation of successive queries can be adjusted to leverage the advantages of Renyi-DP and concentrated DP, considering user inputs. Secondly, it aims to assess how these adjustments and the associated composition methods could enhance the utility of the output. The third objective is to implement these findings into the PostgreSQL Anonymizer, resulting in an improved method for interpreting sequences of SPJA queries, with potentially less overall noise in the data compared to existing approaches at an equivalent epsilon leakage budget.

An article on the integration of these contributions is planed at 18 months. The implementation of these ideas into PostgreSQL Anonymizer is planed at 18 months.

2.3 Integrating Semantic Inference Graph in a DP based approach

The focus is on understanding how public knowledge of a database can lead to privacy breaches through inference. Attackers can exploit partial database knowledge to launch initial attacks and refine them as they gain more information.

The first goal is to identify such attacks based on varying levels of knowledge where, mainly, two scenarios emerge: one where incomplete knowledge impacts privacy, and another where a dynamic approach to data protection is considered. In the former, the aim is to quantify epsilon values ensuring robustness against attacks. In the latter, a tailored privacy approach is proposed, requiring less noise compared to traditional differential privacy methods.

An article describing fine knowledge-aware attacks is planed at 24 months. A second article presenting theoretical results from the dynamic approach to data protection and comparing the utility with classical DP approach is planed at 30 months.

3 Organization

The student will be enrolled at the University of Franche-Comté and will conduct her/his doctoral research at the FEMTO-ST laboratory[1] located at the UFR-ST, 16 route de Gray in Besançon. FEMTO-ST supervisors are Jean-François Couchot and Veronika Sonigo.

Weekly meetings will be held with the supervisors. Travels to INRIA-SACLAY/COMETE[2] are planned to allow the student to work in person with Catuscia Palamidessi.

References

[Clo23] Damien Clochard. Data masking with postgresql anonymizer, 2023. Available at https://dalibo.gitlab.io/postgresql_anonymizer/how-to.handout.pdf.

[Com21] European Comission. Laying down harmonised rules on artificial intelligence (artificial intelligence act) and amending certain union legislative acts. 2021. https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex%3A52021PC0206.

[DFY$^+$23] Wei Dong, Juanru Fang, Ke Yi, Yuchao Tao, and Ashwin Machanavajjhala. R2T: instance-optimal truncation for differentially private query evaluation with foreign keys. *SIGMOD Rec.*, 52(1):115–123, 2023.

[1] https://www.femto-st.fr/en

[2] https://team.inria.fr/Comete/

[DR14] Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Found. Trends Theor. Comput. Sci.*, 9(3-4):211–407, 2014.

[DR16] Cynthia Dwork and Guy N. Rothblum. Concentrated differential privacy. *CoRR*, abs/1603.01887, 2016.

[GHIM19] Chang Ge, Xi He, Ihab F. Ilyas, and Ashwin Machanavajjhala. Apex: Accuracy-aware differentially private data exploration. In Peter A. Boncz, Stefan Manegold, Anastasia Ailamaki, Amol Deshpande, and Tim Kraska, editors, *Proceedings of the 2019 International Conference on Management of Data, SIGMOD Conference 2019, Amsterdam, The Netherlands, June 30 - July 5, 2019*, pages 177–194. ACM, 2019.

[JNS18] Noah M. Johnson, Joseph P. Near, and Dawn Song. Towards practical differential privacy for SQL queries. *Proc. VLDB Endow.*, 11(5):526–539, 2018.

[KTH$^+$19] Ios Kotsogiannis, Yuchao Tao, Xi He, Maryam Fanaeepour, Ashwin Machanavajjhala, Michael Hay, and Gerome Miklau. Privatesql: A differentially private SQL query engine. *Proc. VLDB Endow.*, 12(11):1371–1384, 2019.

[Mir17] Ilya Mironov. Rényi differential privacy. In *30th IEEE Computer Security Foundations Symposium, CSF 2017, Santa Barbara, CA, USA, August 21-25, 2017*, pages 263–275. IEEE Computer Society, 2017.

[TGMR22] Yuchao Tao, Amir Gilad, Ashwin Machanavajjhala, and Sudeepa Roy. Dpxplain: Privately explaining aggregate query answers. *Proc. VLDB Endow.*, 16(1):113–126, 2022.

[WZL$^+$20] Royce J. Wilson, Celia Yuxin Zhang, William Lam, Damien Desfontaines, Daniel Simmons-Marengo, and Bryant Gipson. Differentially private SQL with bounded user contribution. *Proc. Priv. Enhancing Technol.*, 2020(2):230–250, 2020.