
2 Year Postdoctoral or Engineer Position: Differential Privacy in PostgreSQL

INSA Centre Val de Loire is offering an opportunity for either a Postdoctoral Researcher or Engineer to join the Petscraft project-team joint with Inria Saclay. We welcome applications from candidates with diverse profiles and backgrounds, as we believe that both types of expertise can bring valuable contributions to our project. The specific responsibilities, objectives, and scope of the position will be tailored to the selected candidate's qualifications, skills, and experience. We are looking for the successful candidate to start as soon as possible, but we are flexible with the start date to accommodate the right individual's availability.

Host. Petscraft project-team, with a flexible location either at INSA Centre Val de Loire, campus Bourges or Inria Saclay centre, Turing building in Palaiseau.

Contacts. Damien Clochard (damien.clochard@dalibo.com) & Cédric Eichler (cedric.eichler@insa-cvl.fr)

Application. Please contact us if you are interested or need additional details. Application can be sent by email with CV and cover letter.

Context. This position is part of the project "Making PostgreSQL Differentially Private for Transparent AI" funded by the french national research agency ANR. The project involves several French research teams working on data protection, from Inria, universities and engineering schools, as well as Dalibo, company expert in PostgreSQL and developing PostgreSQL Anonymizer [1, 2, 3]. The general objective of the project is to propose, implement within PostgreSQL and evaluate a privacy preserving approach for interpreting SQL queries in the sense of differential privacy (DP). These queries will range from the Select-Project-Join-Aggregation (SPJA) form to the export of releases (DUMP) of a part of the database in order to be able to work on it as if it contained no sensitive data. Specifically, the main objective is to extend the anonymization models already integrated in PostgreSQL Anonymizer (pseudonymization, k-anonymization and addition of noise) to other models verifying DP, as well as DP existing and novel relaxations.

Objectives and roadmap.

During the first year, the successful candidate will tightly collaborate with Dalibo to extend PostgreSQL Anonymizer with the most promising (in terms of utility) approaches [4, 5, 6, 7] to answering SPJA queries that guarantee DP (or similar models, such as local differential privacy). The main lock of this task is to design this DP mechanism as an extension of the SQL grammar in order to qualify a specific query or function with DP metadata.

Possible objectives for the second year are varied and can be tailored to the selected candidate, with a wide range between implementation and theory. They may include more theoretical contributions, e.g. regarding novel DP relaxations or the impact of semantics and schema on DP. In terms of implementation, they could include, for example, the development of demonstrators, validation and evaluation platforms for the novel models proposed in the project.

Profile and appreciated skills. Candidates must hold either a PhD or a master (or equivalent) in Computer Science. The following skills are appreciated:

- Knowledge in privacy & anonymization.
- Knowledge & experience in databases and database management systems, particularly SQL and PostgreSQL.
- Proficiency in programming.

Dalibo will provide training on PostgreSQL and PostgreSQL anonymizer. Courses on anonymization and differential privacy can be proposed at INSA CVL.

Salary & benefits. The salary range is 34000 to 38500 €/year, depending on experience, skills, and qualifications. Benefits:

- Subsidized meals
- Free public transport in Bourges
- Possibility of teleworking (2 days/week)
- Professional equipment available (videoconferencing, loan of computer equipment, etc.)
- Access to vocational training
- Social security coverage
- 7 weeks of paid leave per year

References

- [1] Clochard, D.: Postgresql anonymizer (2022), https://labs.dalibo.com/postgresql_anonymizer
- [2] Clochard, D.: Postgresql anonymizer, canal-u. In: Contraintes réglementaires liées aux bases de données (2022), <https://www.canal-u.tv/133435>
- [3] Clochard, D.: Data Masking with PostgreSQL Anonymizer. Dalibo (2023), https://dalibo.gitlab.io/postgresql_anonymizer/how-to.handout.pdf
- [4] Dong, W., Fang, J., Yi, K., Tao, Y., Machanavajjhala, A.: R2t: Instance-optimal truncation for differentially private query evaluation with foreign keys. In: Proceedings of the 2022 International Conference on Management of Data. p. 759–772. SIGMOD '22, Association for Computing Machinery, New York, NY, USA (2022). <https://doi.org/10.1145/3514221.3517844>, <https://doi.org/10.1145/3514221.3517844>
- [5] Johnson, N.M., Near, J.P., Song, D.X.: Practical differential privacy for SQL queries using elastic sensitivity. CoRR **abs/1706.09479** (2017), <http://arxiv.org/abs/1706.09479>
- [6] Kotsogiannis, I., Tao, Y., He, X., Fanaeepour, M., Machanavajjhala, A., Hay, M., Miklau, G.: Privatesql: a differentially private sql query engine. Proc. VLDB Endow. **12**(11), 1371–1384 (Jul 2019). <https://doi.org/10.14778/3342263.3342274>, <https://doi.org/10.14778/3342263.3342274>
- [7] Wilson, R.J., Zhang, C.Y., Lam, W., Desfontaines, D., Simmons-Marengo, D., Gipson, B.: Differentially private SQL with bounded user contribution. CoRR **abs/1909.01917** (2019), <http://arxiv.org/abs/1909.01917>