



UZH, Dept. of Informatics, Binzmühlestr. 14, CH-8050 Zürich

Thomas Mannhart

Prof. Dr. Michael Böhlen
Professor
Phone +41 44 635 43 33
Fax +41 44 635 68 09
boehlen@ifi.uzh.ch

Zürich, January 8, 2020

Bachelor's Thesis
Datenbanktechnologie

Topic: A General-purpose Range Join Algorithm for PostgreSQL

Joins are crucial and expensive operations in database systems. Current database systems offer efficient evaluation algorithms for equality joins that rely on either partitioning or sort merge techniques. Non-equality joins have been investigated much less and fewer optimization techniques exist.

The goal of this Bachelor thesis is to design, implement and integrate a range join algorithm into the kernel of PostgreSQL. The range join algorithm shall be leveraged for evaluating overlap joins that combine tuples with overlapping ranges.

The work includes the following tasks:

- **T1: Literature study and familiarization with PostgreSQL environment**
Study the relevant literature on temporal normalization [2, 1] and PostgreSQL. Implement a prototype of a range join algorithm as part of the executor of PostgreSQL. Describe the range join algorithm and the relevant parts of the code basis of the prototype in a short document.
- **T2: Implementation of an executor algorithm for range joins**
Integrate the range join algorithm into PostgreSQL together with good cardinality and cost estimates. The outcome shall be a fully functional version of PostgreSQL that can be distributed to and run by others.
- **T3: Empirical performance evaluation of the range join algorithm**
Empirically compare your solution with the temporal normalization of Temporal PostgreSQL [2] and index-based approaches that use ranges in PostgreSQL together with

contains predicates.

- **T4: Patch submission for PostgreSQL**

Prepare and submit a patch of your solution to the PostgreSQL community. Interact with the PostgreSQL community for the preparation and submission of your patch.

Supervisor: Anton Dignös and Michael Böhlen

Start date: 17.2.2020

End date: 17.8.2020

References

- [1] A. Dignös. Temporal unification for database management systems. Master's thesis, Free University of Bozen-Bolzano, 2010.
- [2] A. Dignös, M. H. Böhlen, J. Gamper, and C. S. Jensen. Extending the kernel of a relational DBMS with comprehensive support for sequenced temporal queries. *ACM Trans. Database Syst.*, 41(4):26:1–26:46, 2016.

Department of Informatics, University of Zurich



Prof. Dr. Michael Böhlen