

**Title:** Optimization of Regular Path Queries in Graph Databases

**Name:** Nikolay Yakovets

**Affiliation:** Eindhoven University of Technology

**Abstract:**

Regular path queries offer a powerful navigational mechanism in graph databases. Recently, there has been renewed interest in such queries in the context of the Semantic Web. The extension of SPARQL in version 1.1 with property paths offers a type of regular path query for RDF graph databases. While eminently useful, such queries are difficult to optimize and evaluate efficiently, however. We design and implement a cost-based optimizer we call Waveguide for SPARQL queries with property paths. Waveguide builds a query plan—which we call a waveplan (WP)—which guides the query evaluation. There are numerous choices in the construction of a plan, and a number of optimization methods, so the space of plans for a query can be quite large. Execution costs of plans for the same query can vary by orders of magnitude with the best plan often offering excellent performance. A WP's costs can be estimated, which opens the way to cost-based optimization. We demonstrate that Waveguide properly subsumes existing techniques and that the new plans it adds are relevant. We analyze the effective plan space which is enabled by Waveguide and design an efficient enumerator for it. We implement a prototype of a Waveguide cost-based optimizer on top of an open-source relational RDF store. Finally, we perform a comprehensive performance study of the state of the art for evaluation of SPARQL property paths and demonstrate the significant performance gains that Waveguide offers.

**References:**

Yakovets, Nikolay, Parke Godfrey, and Jarek Gryz. "Query planning for evaluating SPARQL property paths." *Proceedings of the 2016 International Conference on Management of Data*. ACM, 2016.

Yakovets, Nikolay, Parke Godfrey, and Jarek Gryz. "WAVEGUIDE: Evaluating SPARQL Property Path Queries." *EDBT*. 2015.