



Universiteit
Leiden
The Netherlands

Progressive Indexes

Timbó Holanda, P.T.

Citation

Timbó Holanda, P. T. (2021, September 21). *Progressive Indexes*. *SIKS Dissertation Series*. Retrieved from <https://hdl.handle.net/1887/3212937>

Version: Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/3212937>

Note: To cite this publication please use the final published version (if applicable).

Propositions
belonging to the thesis

Progressive Indexes

by Pedro Holanda

Sep. 21, 2021

1. Incremental reorganization must be performance-driven and not workload-driven (Chapter 3).
2. For non-selective queries, tuple reconstruction overhead greatly reduces the utility of conventional fine grained indexes. This is conflicting with existing incremental indexing literature (Chapter 6).
3. Incrementally fitting the data to the workload should be done only in the primary storage layer (Chapter 6).
4. Although Progressive Indexes are designed to optimize interactive workloads, they are not sufficient to ensure interactive query response times (i.e., < 0.5 seconds) (Chapter 3, 4 & 6).
5. Although counter-productive at first sight, combining incremental data reorganization and approximate query processing is a promising method to provide interactive query response times regardless of data size and hardware.
6. Incremental data reorganization is a key design feature of Interactive Database Systems.
7. Contrary to the common opinion in the field, the total execution time is not an interesting metric to optimize when developing solutions that aim to optimize interactive workloads.
8. Although it is generally assumed that adaptive indexes are ideal for exploratory data analysis, they ignore the main metric for interactive workloads, which is ensuring queries are executed under interactive times (i.e., queries are fully executed near real-time).
9. Our publication system must change to the digital era. The current system still follows the format designed to print papers into books, which maximizes text on a single page of paper and does not allow the embedding of supporting materials (e.g., interactive code, videos, GIFs) other than static figures.
10. Conference papers are being rejected for the wrong reasons. Most of the review process follows a "can I kill it quickly" checklist instead of encouraging improvement and discussion. Accepting all papers and openly publishing the reviews, reviewers and affiliations, is a possible direction of change.