To sort or not to sort: the evaluation of R-Tree and B+-Tree in transactional environment with ordered result requirement[1]

P.V. Fedotovsky <pavel.v.fedotovsky@gmail.com>
G.A. Erokhin <george.erokhin@gmail.com>
K.E. Cherednik <kirill.cherednik@math.spbu.ru>
K.K. Smirnov <kirill.k.smirnov@gmail.com>
G.A. Chernishev <chernishev@gmail.com>
SPbU, 11 Universitetskaya nab., St. Petersburg, 199034, Russian Federation

Abstract: In this paper we consider multidimensional indexing with the additional constraint of lexicographical ordering. In order to deal with this problem we discuss two well-known tree data structures: R-tree and B-tree. We study the problem in the transactional environment with read committed isolation level. To evaluate these approaches we had implemented these structures (modified GiST ensures concurrency) and provide extensive experiments.

Keywords: R-Tree; B-Tree; Indexing; Multidimensional Indexing; Lexicographical Order; Experimentation; PostgreSQL.

1. Introduction

In this paper we consider the problem of multidimensional indexing with one additional constraint — the lexicographical ordering of the result set. Effective multidimensional indexing is rather old and well-explored topic, however, one can't say that the problem is solved. New approaches continue to emerge. The addition of the ordering requirement further drives this problem into the domain of research activity.

Effective solutions for the problem of multidimensional indexing are needed for geospatial data, CAD systems, multimedia data and also of use for OLAP data.

There are two main approaches for multidimensional indexing: tree-based and hash-based. The former are R-Tree, KDB tree, Octree, X-Tree and many others. The latter are mainly used for nearest neighbor and similarity query evaluation.

[1] This work is partially supported by Russian Foundation for Basic Research grant 12-07-31050.

We are mainly interested in R-Tree because of its popularity in commercial DBMS systems [1]: PostgreSQL, Oracle, Informix, SQLite and MySQL use this approach. This interest proves, that despite being rather old (more than 25 years), R-Tree still may be called industrial-strength technology. Moreover, until recently R-Tree was the only one method of multidimensional indexing in PostgreSQL2.

This work was inspired by participation in ACM SIGMOD Contest 2012. This problem was provided by the contest organizers, as well as benchmarks and example Berkeley DB-based implementation. Our team participated in this contest and was ranked 5th on public tests3.

The problem is formulated as follows: given an n-dimensional space and queries in transactional environment, what kind of data structure should we use for optimal performance?

In order to solve this problem we implemented a prototype of multidimensional transactional index. This index works within read committed isolation level. Our prototype contains both B^+-Tree and R-Tree built around GiST model.

The contribution of this paper is following:

- The validation of our prototypes by comparison with industrial-strength databases: Berkeley DB and PostgreSQL.
- Experimental study of influence of workload parameters on performance of these two structures. These workload parameters include query window size and others.

The rest of this paper is organized as follows. In the next section we provide detailed specification of the task, describe queries and data. Then, in the section 3 we describe two alternative approaches and survey related works. Section 4 contains overview of our system. In the section 5 we provide evaluations and comparisons with PostgreSQL and Berkeley DB.

2. The Task

The task offered at the contest was to build a multidimensional high-throughput in-memory indexing system. The index should support concurrent access by many threads and work within read committed isolation level. This level is chosen because of its popularity in commercial database systems, for example it is the default isolation level in PostgreSQL. The reason of its popularity is that it offers a good trade-off between correctness and performance. According to the study [2] it can offer 2.5 – 3 times better performance than serializable isolation level. Our index resides in-memory and no crash-recovery component is required.

[2] http://www.postgresql.org/docs/9.2/static/spgist.html
[3] http://web.archive.org/web/20120424201336/http://wwwdb.inf.tu-dresden.de/sigmod2012contest/leaderboard/ accessed: 26/04/2014

2.1. Qu

Ol.eries.

There are several possible types of queries:

- Point queries: insert, update, delete and select.
- Range queries — they select a subset of data and the result should be sorted. This type of query is defined by a conjunction of attribute predicates. The individual predicates may be not only be intervals or points, but also a wildcards.

The distribution of query types is described in the specification.

Another important aspect to consider is the admissible amount of operations per transaction. It is specified, that there are no more than few hundred retrieved points per transaction. In particular, the original task states that no more than 200 points are touched by any transaction. This number is justified by the fact that OLTP transactions are very light-weight. For example, the heaviest transaction in TPC-C reads about 200 records [3].

2.2. Data and Workloads

The task statement specifies several data types: INT(4), INT(8) and VARCHAR(512). However, in this work, we had to drop VARCHAR (see section 5 for details). The key consists of several attributes of these data types. The payload is represented by a sequence of bytes.

The data may come in one of several types of distributions: normal, uniform and Zipf's law (each is applied to coordinate independently). In our tests we used only uniform one.

Duplicate keys are allowed, we refer the reader to the web site for the detailed handling description.

In our experiments we heavily rely upon workloads and benchmark driver provided by organizers. These workloads are essentially synthetic datasets. We don't reuse workloads used during the contest, instead we use the provided framework to define our own.

Thorough task specification can be found here[4].

3. Related Work and Architectural Alternatives

In order to solve this problem two architectural approaches may be used. The first one is to use B^+-Tree and concatenate the values of individual coordinates into the composite key. The B^+-Tree [4] is the balanced data structure, which contains values in the leaf nodes while inner nodes contain pointers and intervals. These intervals define the unique path to the leaf.

[4] http://wwwdb.inf.tu-dresden.de/sigmod2012contest/

The strong points of this approach are:

- The overall simplicity of this data structure and general easiness for implementation.
- The abundance of concurrency control mechanisms for this kind of tree [4].
- It is possible to tune one, a lot of cache-conscious modifications exist.
- No need to sort, because keys are already stored in the right order.

Let's review the last item. Suppose that we have a three dimensional index and a query: (1,2,*). In order to evaluate it, we have to find the first entry with prefix "1|2|" and then sequentially scan the tree until prefix mismatch.

However one can name weak points:

- We have to pack and unpack the keys with each comparison.
- Queries containing interval predicates are harder to process.
- This tree may perform poorly with wildcard queries.

The first one is the minor drawback, its cost may be negligible. However, the second and the third are more formidable ones.

The intervals inside attributes can be processed in the same manner as above, but additional checks are needed. This results in additional complexity of the implementation.

Regarding the third item, consider query (1,*,3). In order to evaluate it, we have to find the key starting with a prefix "1", then we have to iterate through all values which have it. It will require a lot more of comparisons, and what is more important, we will be forced to discard a lot of values in the middle. Consider the following leaf level:

$$\mathbf{1|2|3|, 1|2|4|, 1|2|4|, \ldots, 1|2|4|, 1|3|3|.}$$

In this situation we will need only two values: 1|2|3| and 1|3|3|. But we would be forced to iterate through all these values and discard them. The situation becomes grave when we have wildcard condition in the first attribute: (*,2,3). In this case we have to scan the whole index.

One of the most popular data structures used for indexing multidimensional data is R-tree [5]. According to [6] R-tree is a tree data structure, defined by a pair (m, M) with the following properties:

- Each leaf node (unless it is the root) can host up to M entries, whereas the minimum allowed number of entries is m ≤ M/2. Each entry is of the form (mbr, oid), such that mbr is the MBR that spatially contains the object and oid is the object identifier.
- The number of entries that each internal node can store is again between m ≤ M/2 and M. Each entry is of the form (mbr, p), where p is a pointer to a child of the node and mbr is the MBR that spatially contains the MBRs contained in this child.

- The minimum allowed number of entries in the root node is 2, unless it is a leaf (in this case, it may contain zero or a single entry).
- All leaves of the R-tree are at the same level.

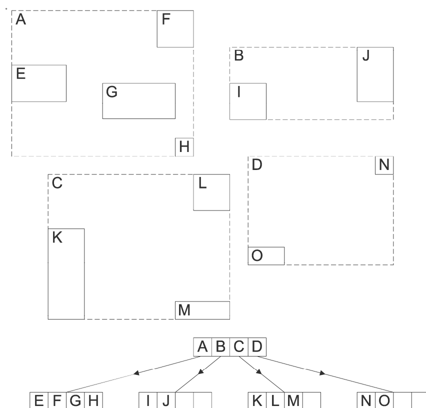

Fig. 1: R-Tree example

Example data and corresponding R-Tree are presented on figure 1. The data consists of boxes marked E-O and boxes A-D denote the MBRs.

R-Tree was proposed first by Antonin Guttman in [7]. This study prompted a wave of research papers and one can say that it gave birth to the new area of research. This research related to development of the new R-Tree variants [1,8,9], niche approaches [6,9], split techniques [10–12], concurrency techniques [13,14] etc. The study [6] states that there are several dozens of R-Tree variants.

One may think of an R-Tree as a generalization of B^+-Tree:

- Data are also kept in the leaves.
- This data structure is also balanced, all leaf nodes are at the same height.
- Inner nodes keep bounding boxes, which can be considered as a generalization of intervals which guide the search in B^+-Tree.

However, there are several differences:

- It might be necessary to explore more than one path from the root to a leaf in order to locate a key. This is the result of the MBR intersection allowance, which leads to more complex search and traversal algorithms than those of B^+-Tree.
- A node split is ambiguous, determining the optimal node split is a computationally hard problem. A “bad” split may lead to a serious degradation of the performance.

- R-Tree does not contain links to sibling leaves for an easy range query execution.

GiST (Generalized Search Tree) [14] is a “template” index structure which supports extensible set of queries and datatypes. This index can be parameterized by a variety of data structures.

Unlike B^+-Tree based one, this approach would require sorting of the results. This is a significant drawback which may negatively impact performance. The goal of this paper is to evaluate, which of these approaches is better. Intuitively one can say that the outcome should depend on the query selectivity.

4. System Overview

Our system follows classical design guidelines and contains several components:

- A tree data structure. Currently implemented as B^+-Tree and R-Tree. R-Tree is based upon GiST [13], a popular template index structure including concurrency control techniques. This model allows to extend with the means of concurrent access almost any tree conforming to certain requirements. This is a widespread approach and it is used, for example, in PostgreSQL.
- Concurrency control. We used mechanism adapted from [14] with locks, latches and Node Sequence Numbers. Also we provided deadlock resolution mechanism. Eventually, we ensure the read committed isolation level. However currently our prototype lacks logging and recovery features.
- Memory manager. It is a well-known fact that a standard memory manager can’t provide optimal performance for the whole range of applications and sometimes it is desirable to find or implement a specifically-tailored one. Our memory manager is essentially a wrapper which intercepts new and delete calls to make use a pool of free blocks.
- Sorting of the results. In order to solve the problem one must present lexicographically sorted results. While B^+-Tree provides already ordered results, R-Tree does not. Our R-Tree implementation sorts the results via merge-sort (we keep sorted data inside boxes).
- Deletion of records. In our implementation we don’t delete records, instead, we mark them as “deleted” and take this into account during the processing.

5. Validation and Experiments

5.1. Validation

We validated our implementation in two ways. First, we used public unit-tests supplied by the contest organizers. These unit-tests ensured correctness of an isolation level (read committed) implementation and several other implementation

issues. We also extended basic set of test cases with new ones. Then, our implementation participated in the contest [15].

5.2. PostgreSQL validation and tuning

We also compared our implementation with PostgreSQL v9.1 database system. This step was needed to check the relative level of achieved performance and general transferability of results. We implemented a simple wrapper application which directed queries to PostgreSQL. PostgreSQL uses a disk-based GiST index, while our prototype is an in-memory one. Also, our prototype lacks a logging and recovery component. Thus, in order to conduct fair tests we had to simulate in-memory index in PostgreSQL.

To completely eliminate slow disk-related operations we placed database cluster on tmpfs. This way we can be sure that every operation PostgreSQL performs (logging, committing, buffers flushing, etc.) does not involve interactions with a hard drive.

Other important implementation aspects included:

- Wrapper connection pooling. We used a pool of connections inside our wrapper to eliminate the cost of connection creation every time a transaction is executed.
- We parameterized GiST with cube data structure.
- To eliminate overheads related to durability we turned off: fsync, full page writes and synchronous commit. Checkpoint segments setting was left intact.
- We were forced to abandon string datatype due to PostgreSQL cube restrictions (only float parameters supported).
- PostgreSQL runs in read committed isolation level by default.

Unfortunately, due to several reasons, we were not able to completely approach the performance of our system. First, unlike BDB, PostgreSQL needs to maintain not only the index, but also a table. Second, calls to PostgreSQL via connections are less effective than the direct function calls. The last issue is the security checks which were also left intact.

5.3. Hardware and software setup

For the first group of experiments (comparison with PostgreSQL and Berkeley DB) used the following hardware and software setup:

- Intel Core i7-2630QM, 2.00 GHz, Hyper-Threading Enabled, L1 Cache 64KB, L2 Cache 256 (per core), L3 Cache 6MB, 6GB RAM
- x86_64 GNU/Linux, kernel 3.5.0-21, gcc 4.7.2
- PostgreSQL 9.1.7

The second group used the more performing one:

- Hardware: 2 x Intel Xeon CPU E5-2660 0 @ 2.20GHz, 64GB RAM, MB S2600GZ
- Software: Linux Ubuntu 3.2.0-29-generic x86_64, GCC 4.6.3

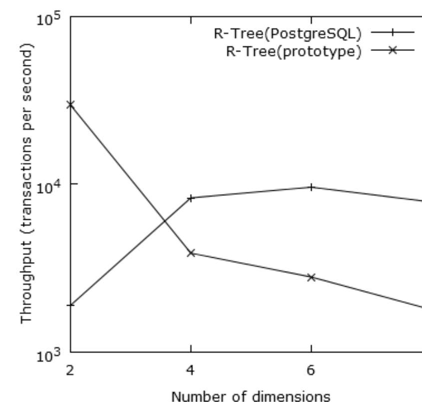

Fig. 2: *Performance of PostgreSQL and our prototype (R-Tree).*

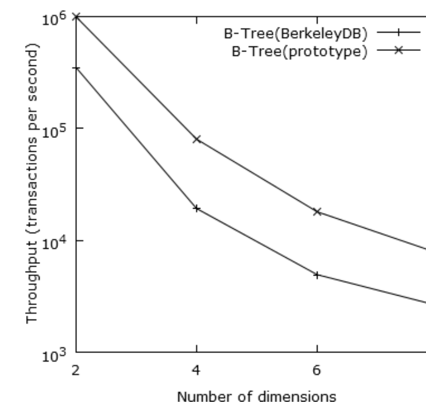

Fig. 3: *Performance of Berkeley DB and our prototype (B+-Tree).*

5.4. Comparison with PostgreSQL and Berkeley DB

In this section we provide a comparison of our prototypes with industrial strength systems. The wrapper for Berkeley DB was provided by the organizers, PostgreSQL wrapper was developed by the authors (its architecture was described earlier). We compare the performance varying the number of dimension and use single 64MB index. The query type distribution is the same as in the original contest task, uniformly distributed data was used. The results are presented on figures 2 and 3.

We can see:

- Our prototypes are comparable to industrial ones in terms of overall performance.
- The solution which uses R-Tree significantly differs from B^+-Tree in terms of performance. This difference has prompted us into further investigation, which resulted in this paper.

5.5. Experimental Evaluation

The goal of this paper is to evaluate, what is better: to use R-Tree and to sort or not to sort with B^+-Tree, but risk excess comparisons.

Table 1: Measured exponent value for R-Tree and B+-Tree, second hardware setup

Tree type	2	4	6	8
R-Tree (64MB)	-0.389 ± 0.052	-0.227 ± 0.011	-0.117 ± 0.012	0.026 ± 0.028
R-Tree (512MB)	-0.547 ± 0.053	-0.250 ± 0.013	-0.135 ± 0.017	-0.071 ± 0.007
B^+-Tree (64MB)	0.482 ± 0.011	0.632 ± 0.025	0.465 ± 0.051	0.333 ± 0.028
B^+-Tree (512MB)	0.476 ± 0.020	0.626 ± 0.038	0.470 ± 0.039	0.321 ± 0.031

In order to solve this problem we had conducted a series of experiments. In these experiments we evaluate the performance of two systems, while varying the query selectivity. We separately consider the following dimensions: 2, 4, 6, 8. We had considered indexes of two sizes: 64 and 512 MB, uniform data distribution. We concentrate on the most interesting query type, which present in the original contest workload: a range query without wildcard predicates. These experiments were conducted using our prototypes, which we had described in the previous section. The reason of this switch is the time it takes to construct an index by PostgreSQL DBMS and also the query plan problem. The plans which are generated by the optimizer are essentially the following: at first, perform index scan (e.g. read all R-Tree boxes), then sort the results. It is impossible to push down sorting in PostgreSQL because its GiST selection method doesn't use merge-sort. This is a critical drawback, because in our task we select at most 200 entries. Thus, our prototype can read only a part of the data and don't sort all the content of the touched boxes. The query plan problem is not an issue in BDB, because of the simplicity of BDB and the fact that B^+-Tree is already sorted.

The graphs describing the experimental results are presented on Figures 4-11.

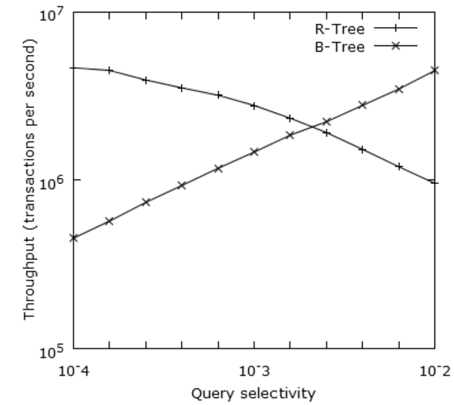

Fig. 4: Performance of R-Tree and B-tree indexes, 64MB, 2 dimensions.

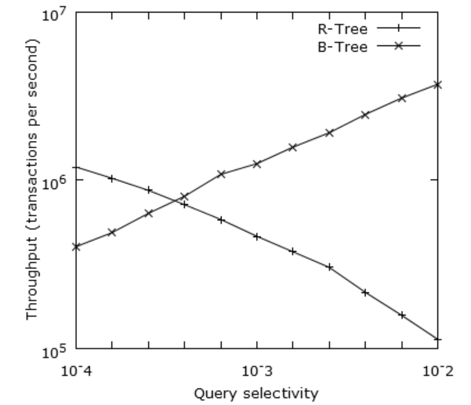

Fig. 5: Performance of R-Tree and B-tree indexes, 512MB, 2 dimensions.

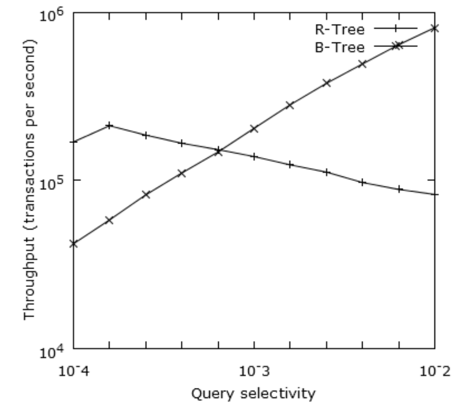

Fig. 6: Performance of R-Tree and B-tree indexes, 64MB, 4 dimensions.

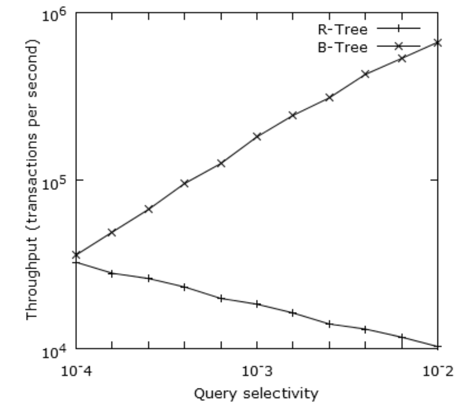

Fig. 7: Performance of R-Tree and B-tree indexes, 512MB, 4 dimensions.

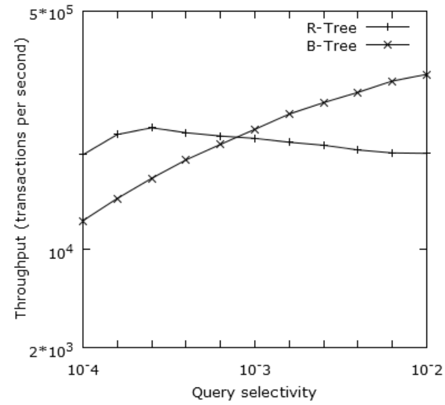

Fig. 8: Performance of R-Tree and B-tree indexes, 64MB, 6 dimensions.

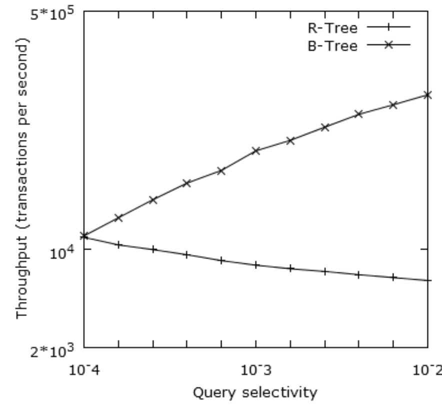

Fig. 9: Performance of R-Tree and B-tree indexes, 512MB, 6 dimensions.

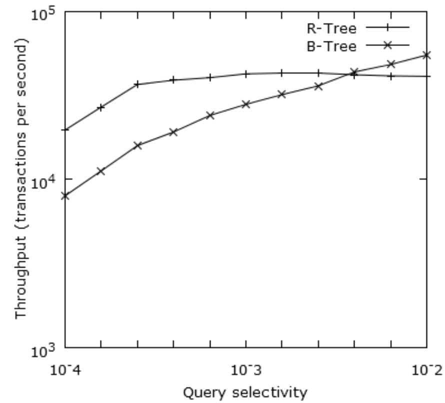

Fig. 10: Performance of R-Tree and B-tree indexes, 64MB, 8 dimensions.

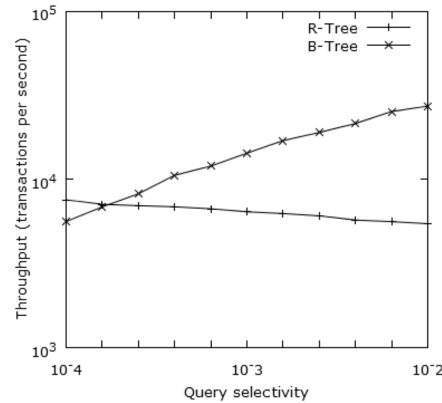

Fig. 11: Performance of R-Tree and B-tree indexes, 512MB, 8 dimensions.

Note the double logarithmic scales, which we used in order to illustrate our finds. They are the following:

- The throughput of the system depends on a query selectivity. This dependence can be described by the power law:
- $P = a * S^b$,
- where P denotes the throughput, S — query selectivity, a and b are parameters. The graphs show this kind of dependency by the straight line. This approximately linear dependency persists in all considered dimension sizes.
- The considered query type affects the performance of the systems in the following way: the performance of R-Tree degrades as the value of query selectivity decreases, while at the same time B^+-Tree performance increases.
- As the number of dimensions increases, the exponent b changes in the way shown in the Table 1. Increasing the dimensionality leads to b decrease in case of R-Tree, i.e. having more dimensions lowers impact of query selectivity. There is no manifested trend in B^+-Tree behaviour.
- There is no simple way to determine intersection point of R-Tree and B-Tree, it depends on number of dimensions and index size.

Table 2: Measured exponent value for R-Tree and B+-Tree, 64MB only, both hardware setups

Tree type	2	4	6	8
R-Tree (1st setup)	-0.369 ± 0.041	-0.220 ± 0.003	-0.115 ± 0.008	0.026 ± 0.032
R-Tree (2nd setup)	-0.389 ± 0.052	-0.227 ± 0.011	-0.117 ± 0.012	0.026 ± 0.028
B^+-Tree (1st setup)	0.478 ± 0.006	0.635 ± 0.034	0.467 ± 0.032	0.332 ± 0.033
B^+-Tree (2nd setup)	0.482 ± 0.011	0.632 ± 0.025	0.465 ± 0.051	0.333 ± 0.028

Despite that the primary goal of this paper is not fully reached, we have obtained a very promising result — the power-law dependency. One of the possible ways of the advancement of this research is to verify, whether the calculated exponent values depend on a hardware setup or not. If we could prove that these value are hardware-independent, then they are fundamental properties of the data structures (at least, implementations) and the sets of queries.

We took the first step the first step of this way and re-evaluated exponent value on different hardware platforms. We had used the first setup (see section 5.3) as this different platform. Unfortunately we could only re-evaluate 64MB workload due to memory limit. The measured values can be found in Table 2 and the corresponding graphs are presented on Figures 12-19. For the easiness of analysis they had been provided with their well-performing counterparts.

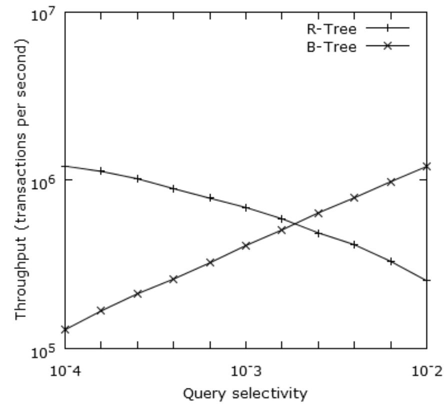

Fig. 12: Performance of R-Tree and B-tree indexes, 64MB, 2 dimensions, 1st configuration.

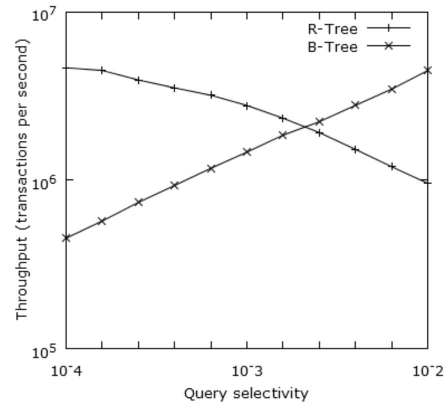

Fig. 13: Performance of R-Tree and B-tree indexes, 64MB, 2 dimensions, 2nd configuration.

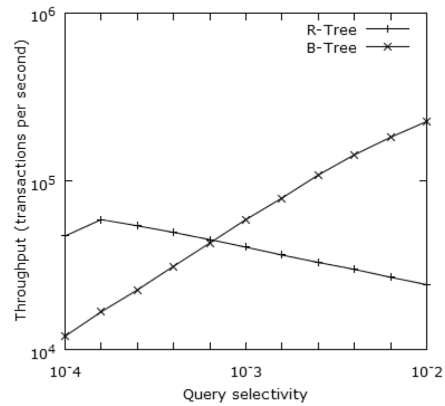

Fig. 14: Performance of R-Tree and B-tree indexes, 64MB, 4 dimensions, 1st configuration.

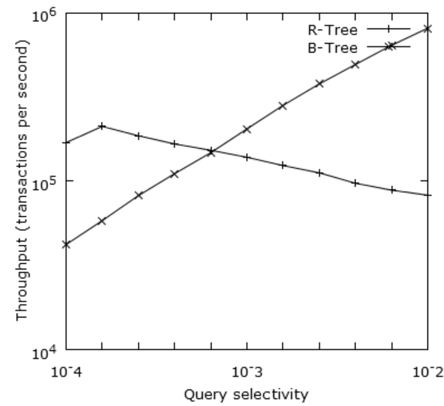

Fig. 15: Performance of R-Tree and B-tree indexes, 64MB, 4 dimensions, 2nd configuration.

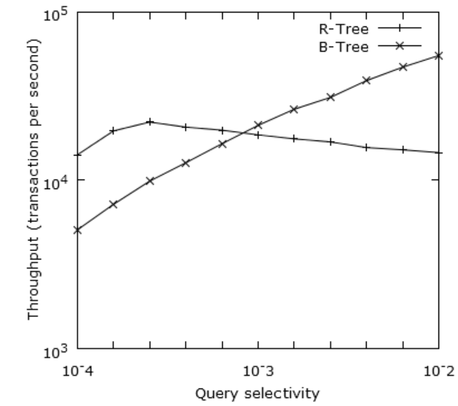

Fig. 16: Performance of R-Tree and B-tree indexes, 64MB, 6 dimensions, 1st configuration.

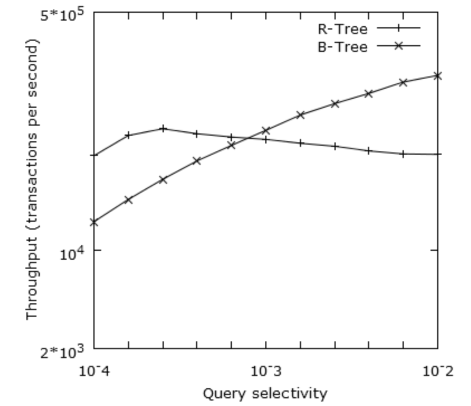

Fig. 17: Performance of R-Tree and B-tree indexes, 64MB, 6 dimensions, 2nd configuration.

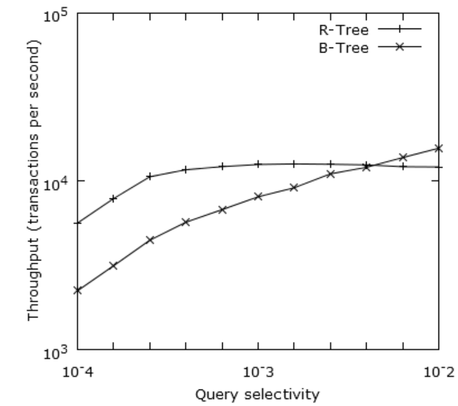

Fig. 18: Performance of R-Tree and B-tree indexes, 64MB, 8 dimensions, 1st configuration.

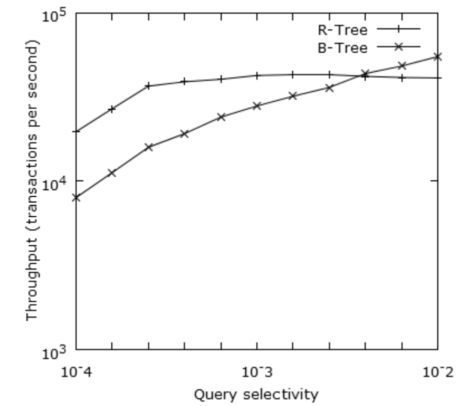

Fig. 19: Performance of R-Tree and B-tree indexes, 64MB, 8 dimensions, 2nd configuration.

As we can see, the values obtained on different hardware setups match each other, so it is very likely that hypothesis of hardware independence holds true and should be carefully examined.

6. Conclusions

In this paper we have considered the problem of multidimensional point indexing under additional restriction: ordering of the result set. We have experimentally evaluated two data structures — R-Tree and B^+-Tree on uniformly distributed data. The experiments allowed us to establish the impact of the query selectivity on system performance as power function. Also we examined the dependency of power-law parameters on dimension cardinality. As a future work we will provide more empirical evidence to the hypothesis of independence of power-law exponent on index size. Recommendation for B-Tree and R-Tree user: unfortunately, we were not able to find an easy way to calculate intersection point, so workloads should be evaluated ad hoc.

7. Acknowledgements

We would like to thank organizers of ACM SIGMOD Programming Contest'12 for providing a benchmark, data generator, unit tests and Berkeley DB wrapper implementation. This work is partially supported by Russian Foundation for Basic Research grant 12-07-31050.

References

[1]. Beckmann N., Seeger B. A revised R*-tree in comparison with related index structures // Proc. of the 2009 ACM SIGMOD International Conference on Management of data, SIGMOD '09. ACM, New York, NY, USA, 2009. P. 799–812.

[2]. Bober P.M., Carey M.J. On mixing queries and transactions via multiversion locking // Proc. of the Eighth International Conference on Data Engineering. IEEE Computer Society, Washington, DC, USA, 1992. P. 535–545.

[3]. Stonebraker M., Madden S., Abadi D.J., Harizopoulos S., Hachem N., Helland P. The end of an architectural era: (it's time for a complete rewrite) // Proc. of the 33rd international conference on Very large databases, VLDB'07. VLDB Endowment, 2007. P. 1150–1160.

[4]. Weikum G., Vossen G. Transactional information systems: theory, algorithms, and the practice of concurrency control and recovery. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2001. 852 P.

[5]. Papadopoulos A.N., Corral A., Nanopoulos A., Theodoridis Y. R-Tree (and Family). In LING LIU and M. TAMER J OZSU, eds. Encyclopedia of Database Systems. Springer US, 2009. P. 2453–2459. doi: 10.1007/978-0-387-39940-9_300.

[6]. Manolopoulos Y., Nanopoulos A., Papadopoulos A.N., Theodoridis Y. R-Trees: Theory and Applications (Advanced Information and Knowledge Processing). Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2005.

[7]. Guttman A. R-trees: a dynamic index structure for spatial searching // SIGMOD Rec. - June 1984. – No. 14(2). P 47–57.

[8]. Kamel I., Faloutsos C. Hilbert Rtree: An Improved R-tree using Fractals // Proc. of the 20th International Conference on Very Large Data Bases, VLDB'94. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1994. P. 500–509.

[9]. Sellis T.K., Roussopoulos N., Faloutsos C. The R+-Tree: A Dynamic Index for Multi-Dimensional Objects // Proc. of the 13th International Conference on Very Large Data

Bases, VLDB'87. Morgan Kaufmann Publishers Inc. San Francisco, CA, USA, 1987. P. 507–518.

[10]. Al-Badarneh A.F., Yaseen Q., Hmeidi I. A new enhancement to the R-tree node splitting // J. Inf. Sci. –feb 2010. –No 36(1). – P 3–18.

[11]. Ang C., Tan T. New linear node splitting algorithm for R-trees // Advances in Spatial Databases Michel Scholl and Agnes Voisard eds vol. 1262 of Lecture Notes in Computer Science. Springer Berlin / Heidelberg, 1997. P. 337–349. doi 10.1007/3-540-63238-7_38.

[12]. Brakatsoulas S., Pfoser D., Theodoridis Y. Revisiting R-Tree Construction Principles // Proc. of the 6th East European Conference on Advances in Databases and Information Systems, ADBIS'02. London, UK, 2002. Springer-Verlag. P 149–162.

[13]. Hellerstein J.M., Naughton J.F., Pfeffer A. Generalized Search Trees for Database Systems // Proc. of the 21th International Conference on Very Large Data Bases, VLDB'95. Morgan Kaufmann Publishers Inc, San Francisco, CA, USA, 1995. P. 562–573.

[14]. Kornacker M., Mohan C., Hellerstein J.M. Concurrency and recovery in generalized search trees // SIGMOD Rec. – June 1997. –No. 26(2). – P 62–72.

[15]. ACM SIGMOD Programming Contest'12. http://wwwdb.inf.tu-dresden.de/sigmod2012contest/. Accessed: 11/09/2012.

Сортировать или нет: экспериментальное сравнение R-Tree и *B*⁺-Tree в транзакционной системе для упорядоченной выдачи

Федотовский П.В. <pavel.v.fedotovsky@gmail.com>
Ерохин Г.А. <george.erokhin@gmail.com>
Чередник К. Е. <kirill.cherednik@math.spbu.ru>
Смирнов К.К. <kirill.k.smirnov@gmail.com>
Чернышев Г.А. <chernishev@gmail.com>
Санкт-Петербургский государственный университет, г. Санкт-Петербург, Россия

Аннотация: В данной работе мы изучаем задачу многомерного индексирования с учетом дополнительного требования – лексикографической упорядоченности результатов запроса. Для решения этой задачи мы рассматриваем две хорошо известные структуры данных – R-дерево и **B+**-дерево, которые используются в транзакционной системе с использованием уровня изоляции read committed. Для сравнения подходов мы реализовали эти структуры (параллельный доступ обеспечивается с помощью GiST) и провели с их помощью ряд экспериментов, результаты которых и представлены в статье.

Ключевые слова: R-дерево; B+-дерево; Индексирование; Многомерное индексирование; Лексикографический порядок; Эксперименты; PostgreSQL.