### MAKE BTREE\_GIST FASTER

### **BERND HELMLE**

www.cybertec-postgresql.com



PROUD CONTRIBUTOR TO



### WHAT'S IT ALL ABOUT?

- GiST (Generalized Search Tree)
- Framework for somehow specialized indexes
- btree\_gist: extension available in PostgreSQL
- Implements btree-like operators/behavior with GiST



### **USAGE EXAMPLE (1)**

nearest-neighbour-search

```
1 CREATE TABLE test (a int4);
2 -- create index
3 CREATE INDEX testidx ON test USING GIST (a);
4
   • • •
5 -- nearest-neighbor search: find the ten entries closest to "42"
6 SELECT *, a <-> 42 AS dist FROM test ORDER BY a <-> 42 LIMIT 10;
```



### USAGE EXAMPLE (2)

Exclusion Constraints

```
1 CREATE TABLE zoo (
2 cage INTEGER,
3 animal TEXT,
4 EXCLUDE USING GIST (cage WITH =, animal WITH <>)
5 );
```



### WE CAN'T HAVE NICE THINGS WITHOUT CAVEATS

- Two build methods for GiST: sorted and buffered
- First one is used, if sortsupport for opclasses used by the index is available
- If not, use buffered: inserts each tuple one by one into the new index
- btree\_gist opclass lacks sortsupport

# by the index is available to the new index



### SOLUTION

- sortsupport API: Implement sortsupport for btree\_gist opclass.
- Extend that idea to builtin range types. . .



## THE PATCH (1)

- Original idea implemented by Andrey Borodin in 2020 (https://commitfest.postgresql.org/patch/2824/)
- Committed in PostgreSQL, but reverted
- Some unresolved problems with correct handling of specific datatypes.





### THE PATCH (2)

- New attempt by Christoph Heiss and me 2022
- Similar problems
- Needed some time to figure out everything



### FINALLY

- commit e4309f73f698851a2f7d49ca5e98e3e188400891 1
- Author: Heikki Linnakangas <heikki.linnakangas@iki.fi> 2
- Date: Thu Apr 3 13:46:35 2025 +0300 3
  - Add support for sorted gist index builds to btree\_gist
- 6 This enables sortsupport in the btree\_gist extension for faster builds 7 of gist indexes. 8
- Sorted gist index build strategy is the new default now. Regression 10 11 tests are unchanged (except for one small change in the 'enum' test to 12 add coverage for enum values added later) and are using the sorted 13 build strategy instead.
- 14

4

5

9

- 15 One version of this was committed a long time ago already, in commit 9f984ba6d2, but it was quickly reverted because of buildfarm 16 failures. The failures were presumably caused by some small bugs, but 17 we never got around to debug and commit it again. This patch was 18 written from scratch, implementing the same idea, with some fragments 19 and ideas from the original patch. 20
- 21
- 22 Author: Bernd Helmle <mailings@oopsware.de>
- Author: Andrey Borodin <x4mmm@yandex-team.ru> 23
- 24 Discussion: https://www.postgresql.org/message-id/64d324ce2a6d535d3f0f3baeeea7b25beff82ce4.camel@oopsware.de



### FINALLY

4

5

6

7

8

9

16

- commit e9e7b66044c9e3dfa76fd1599d5703acd3e4a3f5 1
- Author: Heikki Linnakangas <heikki.linnakangas@iki.fi> 2
- Date: Wed Apr 2 19:51:28 2025 +0300 3
  - Add GiST and btree sortsupport routines for range types
  - For GiST, having a sortsupport function allows building the index using the "sorted build" method, which is much faster.
- For b-tree, the sortsupport routine doesn't give any new 10 11 functionality, but speeds up sorting a tiny bit. The difference is not 12 very significant, about 2% in cursory testing on my laptop, because the range type comparison function has quite a lot of overhead from 13 14 detoasting. In any case, since we have the function for GiST anyway, 15 we might as well register it for the btree opfamily too.
- 17 Author: Bernd Helmle <mailings@oopsware.de>
- Discussion:https://www.postgresql.org/message-id/64d324ce2a6d535d3f0f3baeeea7b25beff82ce4.camel@oopsware.de 18



### SOME NUMBERS (1)

- init.pgbench script for performance testing 1
- 3 BEGIN;

2

4

- DROP TABLE IF EXISTS test\_dataset; 5
- CREATE TABLE test\_dataset(keyid integer not null, id uuid not null, 6
- block range int4range); 7
- CREATE TEMP SEQUENCE testset\_seq; 8
- INSERT INTO test\_dataset SELECT nextval('testset\_seq'), id, block\_range 9
- FROM test ORDER BY random() LIMIT 10000; 10
- CREATE UNIQUE INDEX ON test\_dataset(keyid); 11
- 12

COMMIT; 13



### SOME NUMBERS (2)

- 1 \set keyid random(1, 10000)
- 2 SELECT id, block\_range FROM test\_dataset WHERE keyid = :keyid Ägset
- 3 SELECT id, block\_range FROM test WHERE id = ':id' AND block\_range &&
- 4 ':block\_range';

d = :keyid Ägset ND block\_range &&



### SOME NUMBERS (3)



Figure 1: GiST Build timing





### SOME NUMBERS (4)



Figure 2: GiST Build timing





### BERND HELMLE SENIOR DATABASE ENGINEER

EMAIL bernd.helmle@cybertec.at

**PHONE** +4917680133292



www.cybertec-postgresql.com



@cybertec-postgresql



www.youtube.com/@cybertecpostgresql



