



Carnegie Mellon University

Add/Drop Index

15-721 Final Project Final Presentation

Group 5

Yesheng Ma

Xueyuan Zhao

Jiaqi Zuo

Goals & Progress

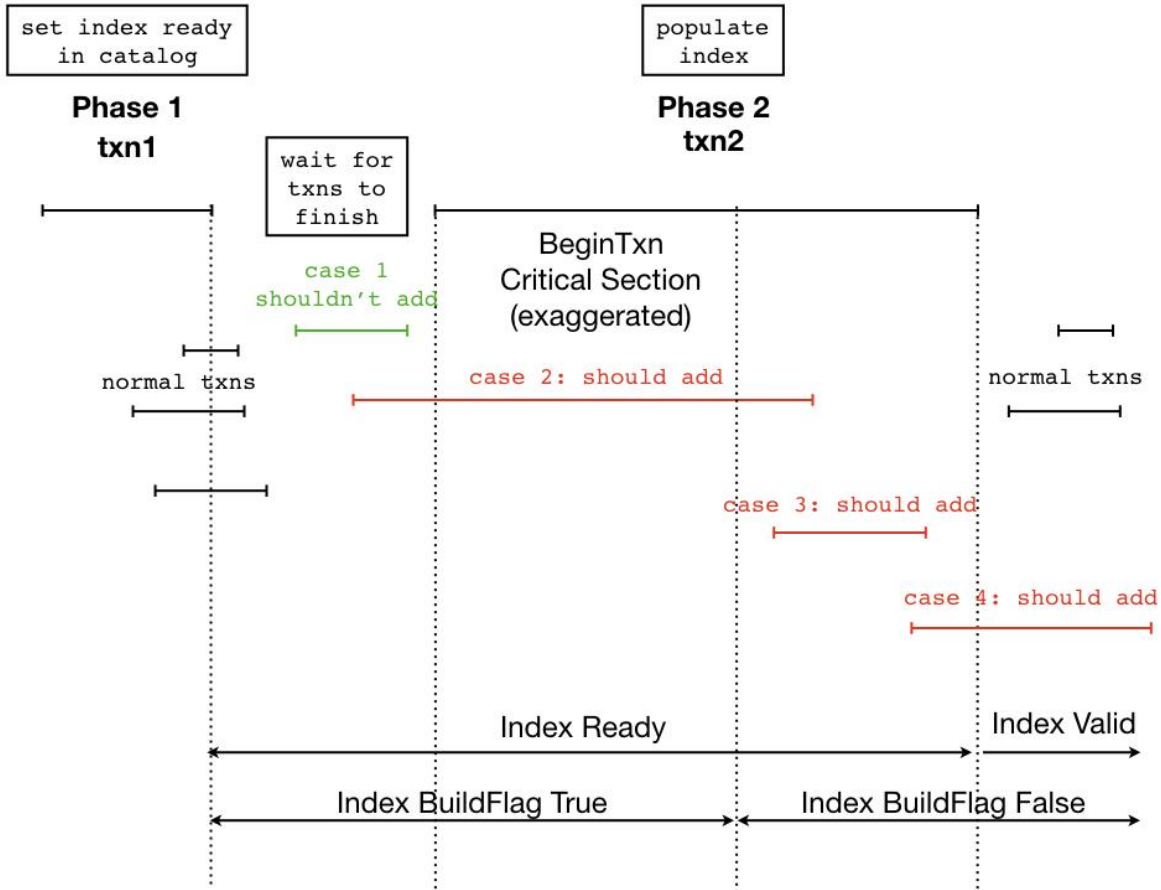
- 75% (DONE): Correctly support add/drop index by blocking all modifications on target table.
- 100% (DONE): Correctly support add/drop index with non-blocking implementation.
- 125% (TODO): Correctly support add/drop index with non-blocking implementation and improve the performance with multi-threading.

Background on MVCC Indexes

- Indexes are associative data structures $\langle \text{Key}, \text{TupleSlot} \rangle$
- Indexes contain *all* keys in a version chain.
- Indexes themselves do not provide visibility info.
- Actual deletions on indexes are only triggered at garbage collection.
- No need to worry about aborted index operations (there are compensate actions for insert/delete).

Design of CREATE INDEX CONCURRENTLY

- Concurrent index creation is super hard to get done right.
- MVCC is *the* burden.
- Some states have to be made visible to other txns right away.
- A two-phase algorithm with all interleavings considered.
- A global IndexBuilder that essentially bypasses the whole MVCC system.
- Critical sections for normal txns are not affected.



Code Quality

- Encapsulate the index related data structure and APIs in the IndexManager class.
- The index building algorithm is well-documented.
- *Core* algorithm is about 500 LoC, relatively easy to verify the correctness.
- But need integration with GC and execution layer to work correctly.

Test & Benchmark

- Unit tests for catalog (pg_index) and index related APIs.
- Basic Add index test: Create an index without concurrent modification.
- End-to-end test: create an index with catalogs involved.
- Interleaving enumeration test (WIP):
 - Handcraft all possible interleavings.
- Performance Benchmarks (TODO):
 - Need to integrate with the GC and execution layer.

Future work

- Scan the table with multiple threads to speed up index building.
- Integrate the execution layer/garbage collection with building index.

Thank you