**Docket #:** S11-171

# Count Indexing for Fast Updates in Column Store Databases

Researchers in Prof. Michael Genesereth's laboratory have developed "count indexes", a unique indexing scheme to efficiently update run-length encoded columns in column stores. This method significantly lowers the update complexity to be logarithmic in the number of the tuples as compared to current methods that update a relation in time that is linear in the number of the tuples. The proposed indexing scheme can be generalized to other compression schemes and supports in-place updates of tuples (both inserts as well as deletes). Count indexing enables the "best of both worlds" scenario in column stores by supporting fast in-place updates while allowing efficient look-ups.

## Applications

- **Enterprise databases** - indexing scheme for updates of column store systems, including run-length and bitmap-encoded compressions
- **Bitmaps and video compression** - indexing scheme for updates and bulk inserts into run-length sequences without decompressing the sequence

## Advantages

- **Fast** - attributes can be updated in time that is logarithmic in the number of runs in the respective columns as opposed to linear
- **Independent of compression scheme** - count indexing can be used with run-length encoding, bit-map encoding, block-oriented storage systems, or uncompressed sequences
- **Supports inserts and deletes**

## Patents

- Published Application: 20130097127
- Issued: 9,720,927 (USA)

## Innovators

- Abhijeet Mohapatra
- Michael Genesereth

## Licensing Contact


### Imelda Oropeza

Senior Licensing Manager, Physcial Sciences

Email