

Extend SkyhookDM programmable object storage with statistics, sort/aggregate and data compaction functions

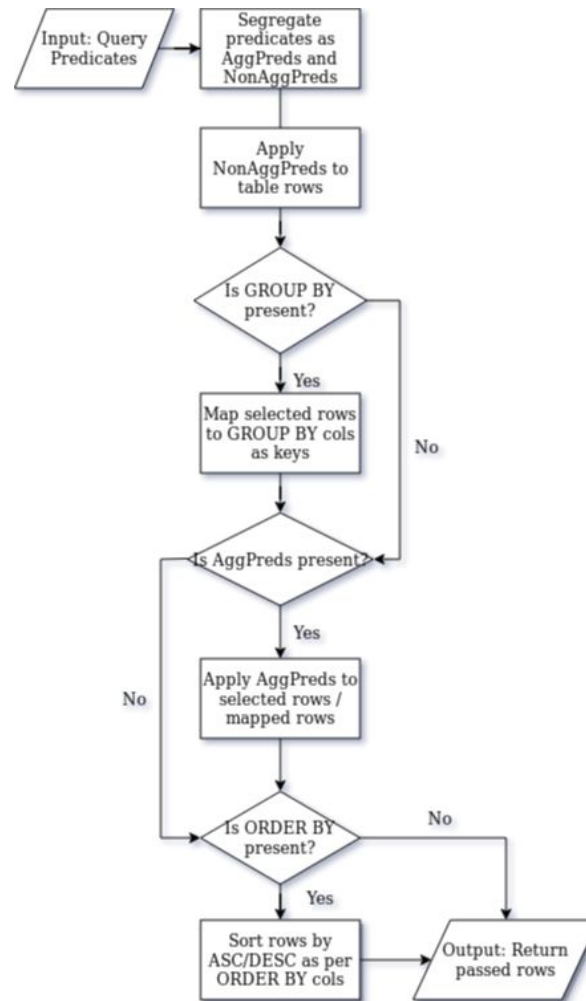
Aditi Gupta

September 03, 2020

Work Done (Phase-1)

- Objective: Extended current aggregation method to include GROUP BY and sort (ORDER BY) for an object's formatted data partitions.
- Implemented using the *split-apply-combine* method inspired by the [pandas](#) library.
- Incorporating support for GROUP BY and ORDER BY has extended Skyhook's capabilities to push down more and more queries to be tackled at object level.

Work Done (Phase-1)



Query Execution Flowchart

Work Done (Phase-2)

- Objective: Implemented a custom method for data statistics collection of an object's formatted data partitions in the form of histograms
- Work done helps in query optimisation by letting the object locally perform access path selection by deciding whether to use indexing or table scans, and also the sequence of operations to be performed depending upon the cardinality of involved tuples.
- Statistics can also be stored locally to let the objects make local decisions as well. ***This is one of the major advantages of programmable object storage in Skyhook.*** This in turn improvises the overall process of physical database design and tuning.

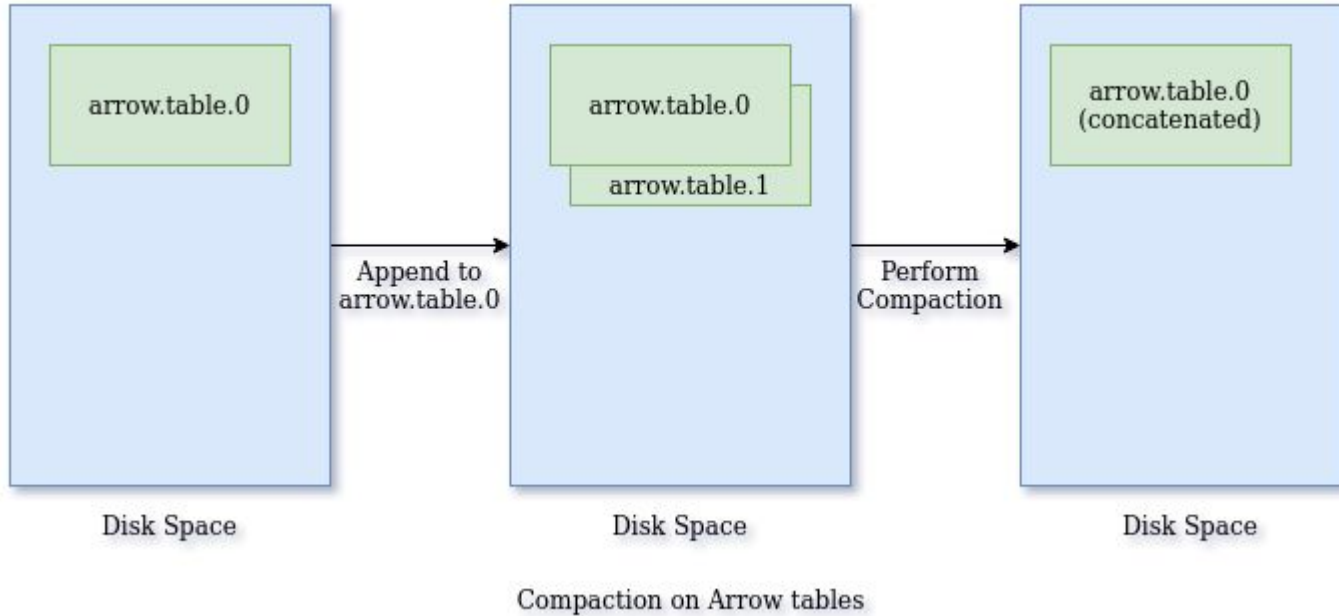
Work Done (Phase-2)

bucket	range	freq	bar
1	[0, 9]	2744	*****
2	[10, 19]	5630	*****
3	[20, 29]	6383	*****
4	[30, 39]	1290	***
5	[40, 49]	369	*
6	[50, 59]	3541	*****
7	[60, 69]	174	
8	[70, 79]	313	*
9	[80, 89]	171	
10	[90, 99]	65	
11	[100, 109]	2363	*****
12	[110, 119]	51	
13	[120, 129]	115	
14	[130, 139]	32	
15	[140, 146]	11	
16	[150, 159]	187	
17	[160, 169]	24	
18	[170, 177]	33	
19	[180, 189]	19	
20	[191, 199]	24	
21	[200, 200]	795	**

Work Done (Phase-3)

- Objective: Implement compaction of multiple formatted sub-partitions within an object into a single partition.
- Developed an object class method that merges formatted data partitions within an object. Self-contained partitions are written (appended) to objects and over time objects may contain a sequence of independent formatted data structures.
- The compaction operation consolidates the non-contiguous table fragments stored in individual memory blocks for effective memory management. The implemented function merges sequences of independent Arrow tables into a single larger Arrow table and writes a clean partition back into the disk.

Work Done (Phase-3)



Further Details

Please refer to the final report:

<https://github.com/uccross/skyhookdm-ceph/wiki/Google-Summer-of-Code-2020-Report>

For clarifications, get in touch with me at guptaaditi1709@gmail.com