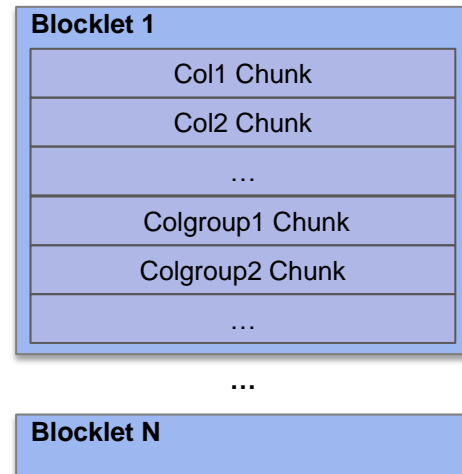


# CarbonData File Format

# Carbon File Structure

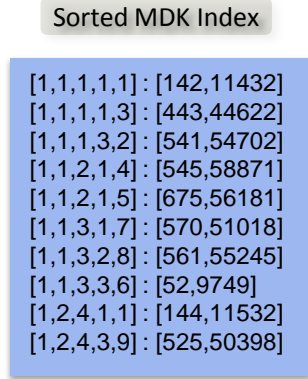
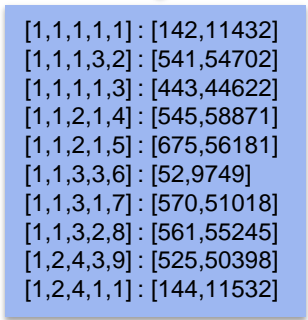
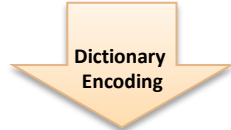
- **Blocklet : A set of rows in columnar format**
  - Default blocklet size: ~120k rows
  - Minimum size for predicate filtering
  - Large size for efficient reading and compression
- **Column chunk : Data for one column/column group in a Blocklet**
  - Column data stored as sorted index
  - Allow multiple columns forms a column group & stored as row-based
- **Footer : Metadata information**
  - File level statistics
  - Schema
  - Blocklet Index & Metadata



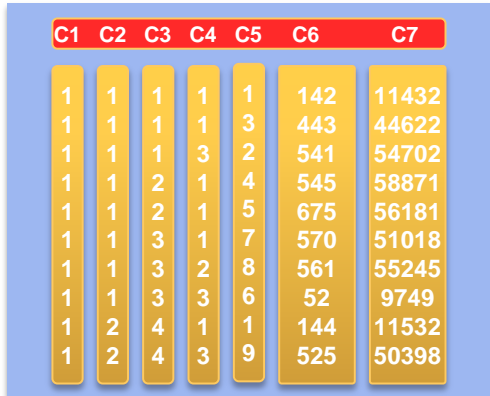
# Blocklet

Column data are sorted along multi-dimensional key

Years	Quarters	Months	Territory	Country	Quantity	Sales
2003	QTR1	Jan	EMEA	Germany	142	11,432
2003	QTR1	Jan	APAC	China	541	54,702
2003	QTR1	Jan	EMEA	Spain	443	44,622
2003	QTR1	Feb	EMEA	Denmark	545	58,871
2003	QTR1	Feb	EMEA	Italy	675	56,181
2003	QTR1	Mar	APAC	India	52	9,749
2003	QTR1	Mar	EMEA	UK	570	51,018
2003	QTR1	Mar	Japan	Japan	561	55,245
2003	QTR2	Apr	APAC	Australia	525	50,398
2003	QTR2	Apr	EMEA	Germany	144	11,532



Blocklet Logical View



# Column Chunk Inverted Index

Configurable inverted index building for column chunk for better compression & binary search within column chunk

Columnar MDK Index

```
[1|1] :[1|1] :[1|1] :[1|1] :[1|1] : [142]:[11432]
[1|2] :[1|2] :[1|2] :[1|2] :[1|10] : [541]:[54702]
[1|3] :[1|3] :[1|3] :[1|3] :[2|2] : [443]:[44622]
[1|4] :[1|4] :[2|4] :[1|4] :[3|3] : [545]:[58871]
[1|5] :[1|5] :[2|5] :[1|5] :[4|4] : [675]:[56181]
[1|6] :[1|6] :[3|6] :[1|6] :[5|5] : [52]:[9749]
[1|7] :[1|7] :[3|7] :[1|7] :[6|6] : [570]:[51018]
[1|8] :[1|8] :[3|8] :[1|10] :[7|7] : [561]:[55245]
[1|9] :[2|9] :[4|9] :[2|8] :[8|8] : [525]:[50398]
[1|10]:[2|10]:[4|10]:[3|9] :[9|9] : [144]:[11532]
```



Run Length Encoding & Compression

Block Level inverted Index

Blocklet Physical View

C1		C2		C3		C4		C5			
d	r	d	r	d	r	d	r	d	r		
1	1	1	1	1	1	1	1	1	1	142	11432
10	10	8	10	3	10	8	7	2	10	443	44622
		2		2		2	10	2		541	54702
		2		2		1	1	1		545	58871
		3		3		3	8	3		675	56181
		3		3		1	1	1		570	51018
		4		2		1	9	4		561	55245
						1		1		52	9749
								5		144	11532
								1		525	50398
								...			



Columnar Store

Blocklet Rows

Dim1 Block	Dim2 Block	Dim3 Block	Dim4 Block	Dim5 Block	Measure1 Block	Measure2 Block
{{1(1-10) (1->10)}}	:[{{1(1-8) (1->8)}}	:[{{1(1-3) (1->3)}}	:[{{1(1-8) (1-7,10)}}	:[{{1(1-2) (1,10)}}	[142]:	[11432]
	[2(9-10) (9->10)}}	[2(4-5) 4->5]	[2 8]	[2 2]	[541]:	[54702]
		[3(6-8) (6->8)}}	[3 9]}	[3 3]	[443]:	[44622]
		[4(9-10) (9->10)}}		[4 4]	[545]:	[58871]
				[5 5]	[675]:	[56181]
				[6 6]	[52]:	[9749]
				[7 7]	[570]:	[51018]
				[8 8]	[561]:	[55245]
				[9 9]}	[525]:	[50398]
					[144]:	[11532]

# Column Group

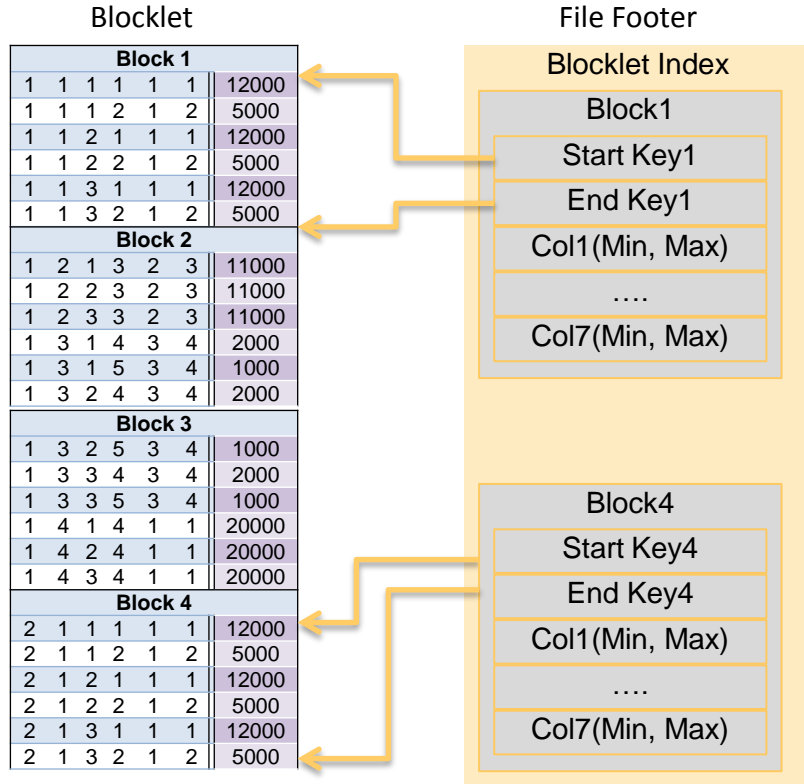
Allow multiple columns form a column group

- stored as a single column chunk in row-based format
- suitable to set of columns frequently fetched together
- saving stitching cost for reconstructing row

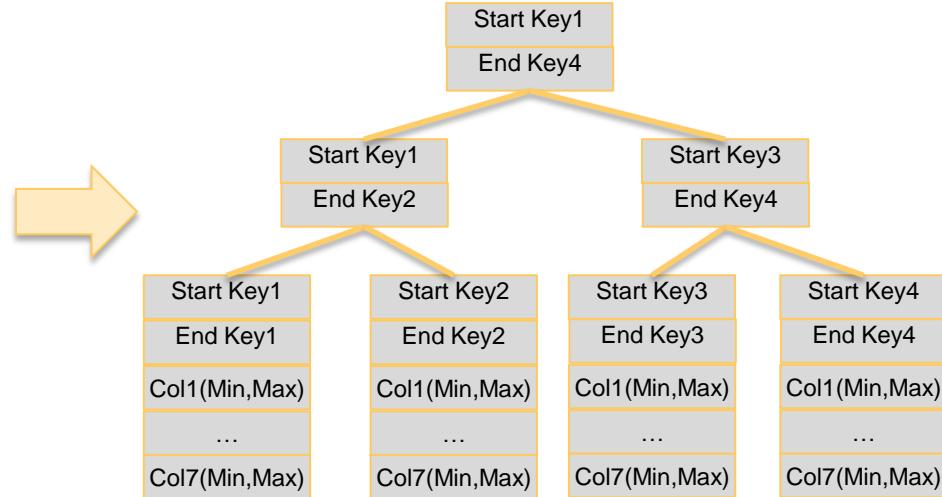
HDFS Block

	Blocklet 1					
	C1	C2	C3	C4	C5	C6
	Col Chunk	Col Chunk	Col Chunk	Col Group Chunk	Col Chunk	Col Chunk
8	10(1)	2(1)	23(3)			
9	10(2)	2(2)	50(1)			
6	10(3)	3(3)	51(2)			
7	10(1)	3(1)	24(1)			
5	10(2)	3(2)	25(2)			
10	10(3)	3(3)	26(3)			
2	10(1)	1(2)	0(1)			
12	11(2)	1(3)	1(2)			
11	11(3)	4(1)	50(3)			
3	11(1)	5(1)	5(1)			
1	11(2)	6(3)	45(3)			
4	18(3)	9(2)	76(2)			

# Blocklet Index



Build in-memory file level MDK index tree for filtering Blocklet



# Encoding & Compression

- Supported Encoding Scheme:
  - DELTA Encoding:
  - Run Length Encoding:
  - BIT\_PACKED:
  - Dictionary Encoding:
    - low cardinality: global dictionary (table level)
      - speedup aggregation
      - less run-time memory usage
      - Lazy decoding
      - Fast distinct counts
    - medium high cardinality: local dictionary (file level)
  - CUSTOM
- Compression Scheme: Snappy

Thank you