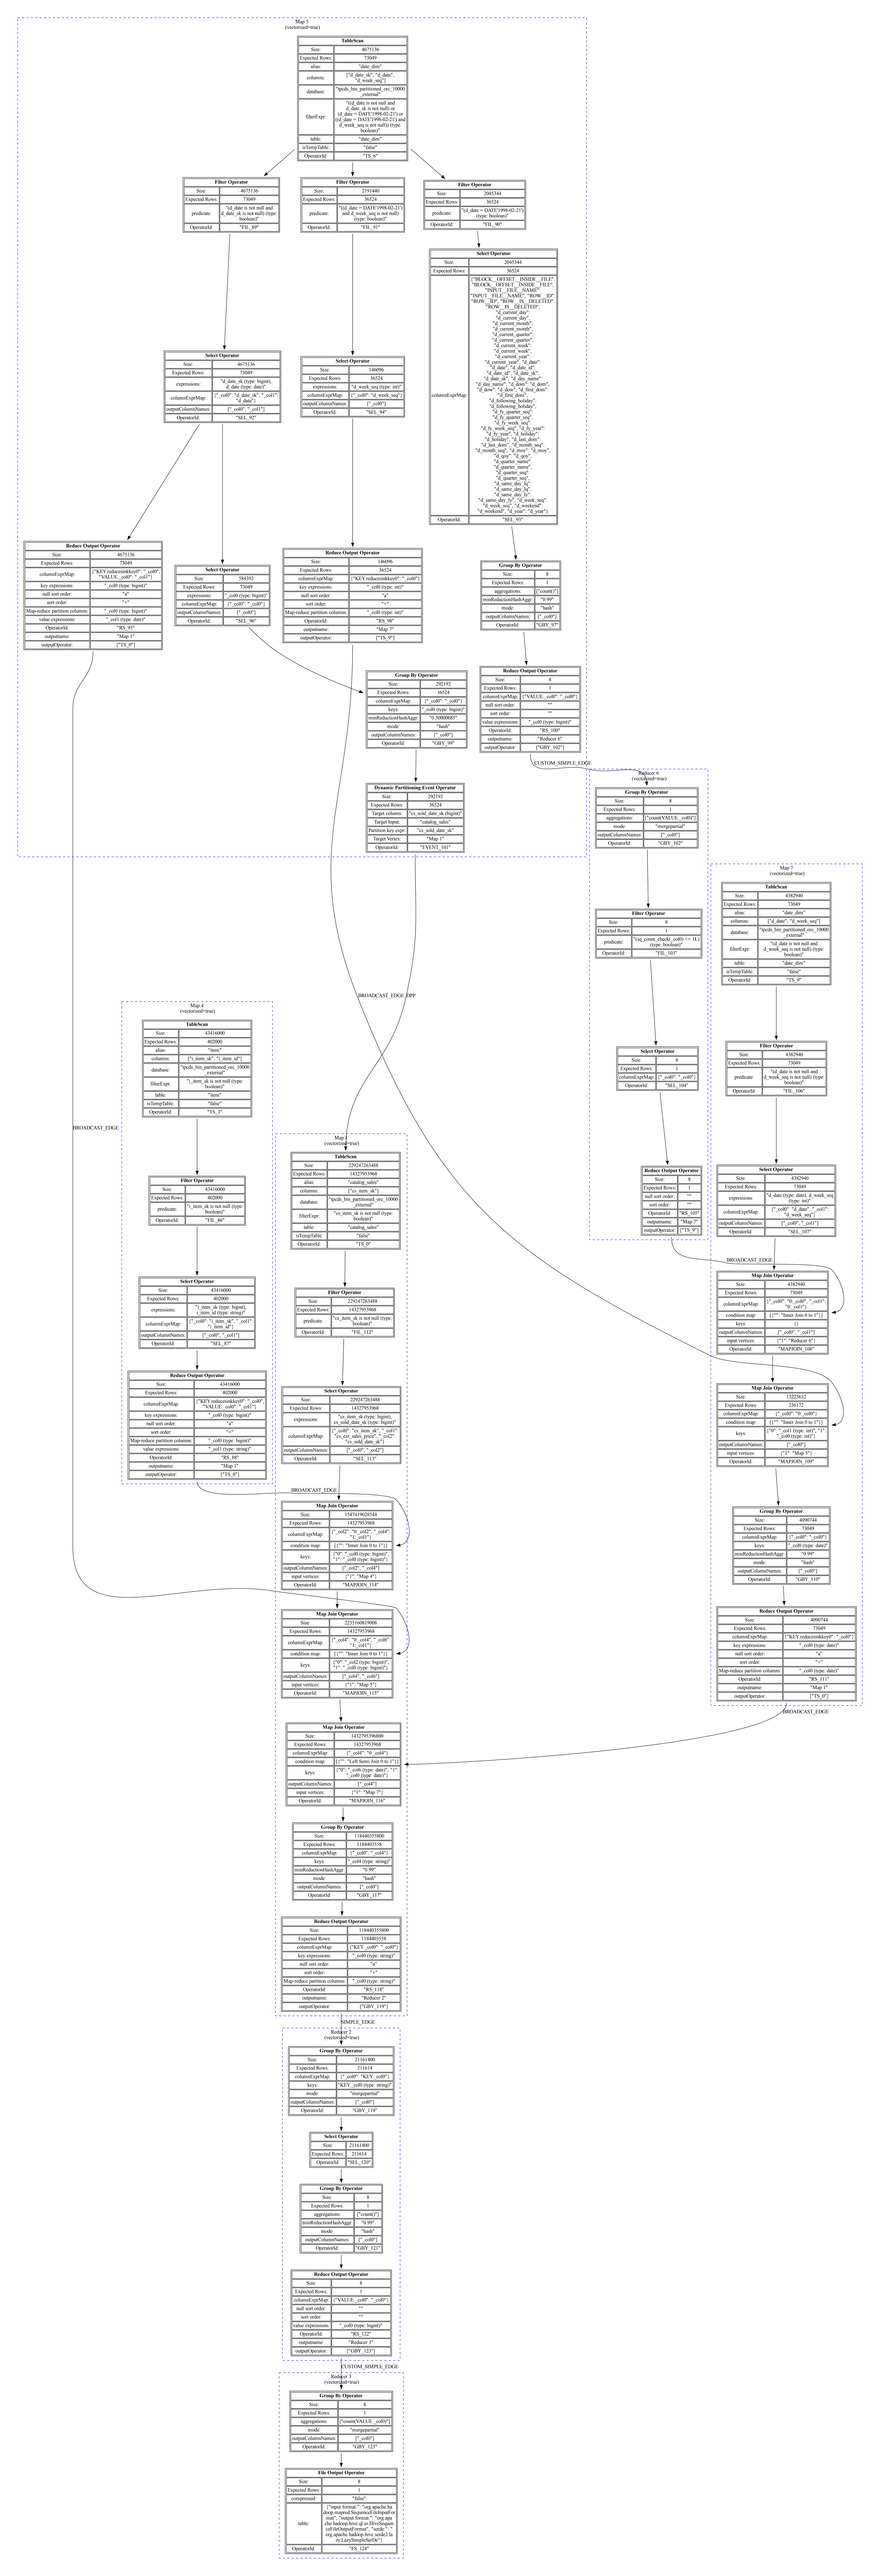
Map 2
(vectorized=true)
TableScan
Size: 4675136
Expected Rows: 73049
alias: "date_dim"
columns: ["d_date_sk", "d_date", "d_week_seq"]
database: "tpcds_bin_partitioned_orc_10000_external"
filterExpr: "((d_date is not null and d_date_sk is not null) or (d_date = DATE'1998-02-21') or (d_date = DATE'1998-02-21') and d_week_seq is not null) (type: boolean)"
table: "date_dim"
isTempTable: "false"
OperatorId: "TS_0"
Filter Operator
Size: 4675136
Expected Rows: 73049
predicate: "(d_date is not null and d_date_sk is not null) (type: boolean)"
OperatorId: "FIL_89"
Filter Operator
Size: 2191440
Expected Rows: 36524
predicate: "((d_date = DATE'1998-02-21') and d_week_seq is not null) (type: boolean)"
OperatorId: "FIL_91"
Filter Operator
Size: 2043344
Expected Rows: 36524
predicate: "(d_date = DATE'1998-02-21') (type: boolean)"
OperatorId: "FIL_90"
Select Operator
Size: 2043344
Expected Rows: 36524
columnExprMap: {"BLOCK__OFFSET__INSIDE__FILE}, "BLOCK__OFFSET__INSIDE__FILE", "INPUT__FILE__NAME}, "INPUT__FILE__NAME", "ROW__ID}, "ROW__ID", "PARTITION__SPEC__ID}, "PARTITION__SPEC__ID", "ROW__ID__FETCH}, "ROW__ID__FETCH", "d_current_day", "d_current_day", "d_current_month", "d_current_month", "d_current_quarter", "d_current_quarter", "d_current_week", "d_current_week", "d_current_year", "d_current_year", "d_date", "d_date", "d_date_id", "d_date_id", "d_date_sk", "d_date_sk", "d_day_name", "d_day_name", "d_dom", "d_dom", "d_dow", "d_dow", "d_first_dom", "d_first_dom", "d_following_holiday", "d_following_holiday", "d_fom", "d_fom", "d_fy_quarter_seq", "d_fy_quarter_seq", "d_fy_week_seq", "d_fy_week_seq", "d_fy_year", "d_fy_year", "d_holiday", "d_holiday", "d_last_dom", "d_last_dom", "d_last_dom", "d_last_dom", "d_month_seq", "d_month_seq", "d_moy", "d_moy", "d_qoy", "d_qoy", "d_quarter_name", "d_quarter_name", "d_quarter_seq", "d_quarter_seq", "d_same_day_lq", "d_same_day_lq", "d_same_day_ly", "d_same_day_ly", "d_time_dim_id", "d_week_seq", "d_week_seq", "d_weekend", "d_weekend", "d_year", "d_year"]
OperatorId: "SEL_93"
Select Operator
Size: 4675136
Expected Rows: 73049
expressions: "d_date_sk (type: bigint), d_date (type: date)"
columnExprMap: {"_col0", "d_date_sk", "_col1", "d_date"}
outputColumnNames: ["_col0", "_col1"]
OperatorId: "SEL_92"
Select Operator
Size: 186096
Expected Rows: 36524
expressions: "d_week_seq (type: int)"
columnExprMap: {"_col0", "d_week_seq"}
outputColumnNames: ["_col0"]
OperatorId: "SEL_94"
Reduce Output Operator
Size: 4675136
Expected Rows: 73049
columnExprMap: {"KEY.reducesinkkey0", "_col0", "VALUE._col0", "_col1"}
key expressions: "_col0 (type: bigint)"
null sort order: "a"
sort order: "+"
Map-reduce partition columns: "_col0 (type: bigint)"
value expressions: "_col0 (type: date)"
OperatorId: "RS_95"
outputname: "Map 1"
outputOperator: ["FS_0"]
Select Operator
Size: 584392
Expected Rows: 73049
expressions: "_col0 (type: bigint)"
columnExprMap: {"_col0", "_col0"}
outputColumnNames: ["_col0"]
OperatorId: "SEL_96"
Reduce Output Operator
Size: 186096
Expected Rows: 36524
columnExprMap: {"KEY.reducesinkkey0", "_col0"}
key expressions: "_col0 (type: int)"
null sort order: "a"
sort order: "+"
Map-reduce partition columns: "_col0 (type: int)"
OperatorId: "RS_98"
outputname: "Map 3"
outputOperator: ["FS_9"]
Group By Operator
Size: 8
Expected Rows: 1
aggregations: ["count()"]
mode: "hash"
outputColumnNames: ["_col0"]
OperatorId: "GBY_97"
Reduce Output Operator
Size: 8
Expected Rows: 1
columnExprMap: {"VALUE._col0"}
null sort order: ""
sort order: ""
value expressions: "_col0 (type: bigint)"
OperatorId: "RS_100"
outputname: "Reducer 6"
outputOperator: ["GBY_102"]
Group By Operator
Size: 292192
Expected Rows: 36524
columnExprMap: {"_col0", "_col0"}
keys: "_col0 (type: bigint)"
minReductionHashAggr: "0.50000685"
mode: "hash"
outputColumnNames: ["_col0"]
OperatorId: "GBY_99"
Dynamic Partitioning Event Operator
Size: 292192
Expected Rows: 36524
Target column: "ss_sold_date_sk (bigint)"
Target Input: "catalog_sales"
Partition key expr: "ss_sold_date_sk"
Target Vertex: "Map 1"
OperatorId: "EVENT_101"
Reducer 6
(vectorized=true)
Group By Operator
Size: 8
Expected Rows: 1
aggregations: ["count(VALUE._col0)"]
mode: "mergepartial"
outputColumnNames: ["_col0"]
OperatorId: "GBY_102"
Filter Operator
Size: 8
Expected Rows: 1
predicate: "(sq_count_check(_col0) <= 1) (type: boolean)"
OperatorId: "FIL_103"
Select Operator
Size: 8
Expected Rows: 1
columnExprMap: {"_col0", "_col0"}
OperatorId: "SEL_104"
Reduce Output Operator
Size: 8
Expected Rows: 1
null sort order: ""
sort order: ""
OperatorId: "RS_105"
outputname: "Map 7"
outputOperator: ["FS_9"]
Map 7
(vectorized=true)
TableScan
Size: 4382940
Expected Rows: 73049
alias: "date_dim"
columns: ["d_date", "d_week_seq"]
database: "tpcds_bin_partitioned_orc_10000_external"
filterExpr: "(d_date is not null and d_week_seq is not null) (type: boolean)"
table: "date_dim"
isTempTable: "false"
OperatorId: "TS_9"
Filter Operator
Size: 4382940
Expected Rows: 73049
predicate: "(d_date is not null and d_week_seq is not null) (type: boolean)"
OperatorId: "FIL_106"
Select Operator
Size: 4382940
Expected Rows: 73049
expressions: "d_date (type: date), d_week_seq (type: int)"
columnExprMap: {"_col0", "d_date", "_col1", "d_week_seq"}
outputColumnNames: ["_col0", "_col1"]
OperatorId: "SEL_107"
BROADCAST_EDGE
Map Join Operator
Size: 4382940
Expected Rows: 73049
columnExprMap: {"_col0", "_col0", "_col1", "_col1"}
condition map: [{"0": "Inner Join 0 to 1"}]
keys: {}
outputColumnNames: ["_col0", "_col1"]
input vertices: {"1": "Reducer 6"}
OperatorId: "MAPJOIN_108"
Map Join Operator
Size: 13235632
Expected Rows: 236172
columnExprMap: {"_col0", "_col0", "_col1", "_col1"}
condition map: [{"0": "Inner Join 0 to 1"}]
keys: {"0": "_col1 (type: int)", "1": "_col0 (type: int)"}
outputColumnNames: ["_col0", "_col1"]
input vertices: {"1": "Map 9"}
OperatorId: "MAPJOIN_109"
Group By Operator
Size: 4090744
Expected Rows: 73049
columnExprMap: {"_col0", "_col0", "_col1", "_col1"}
keys: "_col0 (type: date)"
minReductionHashAggr: "0.99"
mode: "hash"
outputColumnNames: ["_col0", "_col1"]
OperatorId: "GBY_110"
Reduce Output Operator
Size: 4090744
Expected Rows: 73049
columnExprMap: {"KEY.reducesinkkey0", "_col0", "_col1"}
key expressions: "_col0 (type: date)"
null sort order: "a"
sort order: "+"
Map-reduce partition columns: "_col0 (type: date)"
OperatorId: "RS_111"
outputname: "Map 1"
outputOperator: ["FS_0"]
BROADCAST_EDGE
Map 4
(vectorized=true)
TableScan
Size: 43410400
Expected Rows: 402000
alias: "item"
columns: ["i_item_sk", "i_item_id"]
database: "tpcds_bin_partitioned_orc_10000_external"
filterExpr: "i_item_sk is not null (type: boolean)"
table: "item"
isTempTable: "false"
OperatorId: "TS_3"
Filter Operator
Size: 43410400
Expected Rows: 402000
predicate: "i_item_sk is not null (type: boolean)"
OperatorId: "FIL_86"
Select Operator
Size: 43410400
Expected Rows: 402000
expressions: "i_item_sk (type: bigint), i_item_id (type: string)"
columnExprMap: {"_col0", "i_item_sk", "_col1", "i_item_id"}
outputColumnNames: ["_col0", "_col1"]
OperatorId: "SEL_87"
Reduce Output Operator
Size: 43410400
Expected Rows: 402000
columnExprMap: {"KEY.reducesinkkey0", "_col0", "VALUE._col0", "_col1"}
key expressions: "_col0 (type: bigint)"
null sort order: "a"
sort order: "+"
Map-reduce partition columns: "_col0 (type: bigint)"
value expressions: "_col1 (type: string)"
OperatorId: "RS_88"
outputname: "Map 1"
outputOperator: ["FS_0"]
BROADCAST_EDGE
Map 1
(vectorized=true)
TableScan
Size: 229247263488
Expected Rows: 143279539968
alias: "cs_item_sk"
database: "tpcds_bin_partitioned_orc_10000_external"
filterExpr: "cs_item_sk is not null (type: boolean)"
table: "catalog_sales"
isTempTable: "false"
OperatorId: "TS_0"
Filter Operator
Size: 229247263488
Expected Rows: 143279539968
predicate: "cs_item_sk is not null (type: boolean)"
OperatorId: "FIL_112"
Select Operator
Size: 229247263488
Expected Rows: 143279539968
expressions: "cs_item_sk (type: bigint), cs_sold_date_sk (type: bigint)"
columnExprMap: {"_col0", "cs_item_sk", "_col1", "cs_sold_date_sk"}
outputColumnNames: ["_col0", "_col1"]
OperatorId: "SEL_113"
BROADCAST_EDGE
Map Join Operator
Size: 154741902634
Expected Rows: 143279539968
columnExprMap: {"_col0", "0", "_col0", "_col0", "1", "_col1"}
condition map: [{"0": "Inner Join 0 to 1"}]
keys: {"0": "_col0 (type: bigint)", "1": "_col0 (type: bigint)"}
outputColumnNames: ["_col0", "_col1"]
input vertices: {"1": "Map 4"}
OperatorId: "MAPJOIN_114"
Map Join Operator
Size: 223310819808
Expected Rows: 143279539968
columnExprMap: {"_col0", "0", "_col0", "_col0", "1", "_col1"}
condition map: [{"0": "Inner Join 0 to 1"}]
keys: {"0": "_col2 (type: bigint)", "1": "_col0 (type: bigint)"}
outputColumnNames: ["_col0", "_col1"]
input vertices: {"1": "Map 5"}
OperatorId: "MAPJOIN_115"
Map Join Operator
Size: 143279539680
Expected Rows: 143279539968
columnExprMap: {"_col0", "0", "_col0"}
condition map: [{"0": "Left Semi Join 0 to 1"}]
keys: {"0": "_col0 (type: date)", "1": "_col0 (type: date)"}
outputColumnNames: ["_col0"]
input vertices: {"1": "Map 7"}
OperatorId: "MAPJOIN_116"
Group By Operator
Size: 11844035000
Expected Rows: 1184403518
columnExprMap: {"_col0", "_col0"}
keys: "_col0 (type: string)"
minReductionHashAggr: "0.99"
mode: "hash"
outputColumnNames: ["_col0"]
OperatorId: "GBY_117"
Reduce Output Operator
Size: 11844035000
Expected Rows: 1184403518
columnExprMap: {"KEY._col0", "_col0"}
key expressions: "_col0 (type: string)"
null sort order: "a"
sort order: "+"
Map-reduce partition columns: "_col0 (type: string)"
OperatorId: "RS_118"
outputname: "Reducer 2"
outputOperator: ["GBY_119"]
SIMPLE_EDGE
Reducer 2
(vectorized=true)
Group By Operator
Size: 21161400
Expected Rows: 211614
columnExprMap: {"_col0", "KEY._col0"}
keys: "KEY._col0 (type: string)"
mode: "mergepartial"
outputColumnNames: ["_col0"]
OperatorId: "GBY_119"
Select Operator
Size: 21161400
Expected Rows: 211614
OperatorId: "SEL_120"
Group By Operator
Size: 8
Expected Rows: 1
aggregations: ["count(1)"]
minReductionHashAggr: "0.99"
mode: "hash"
outputColumnNames: ["_col0"]
OperatorId: "GBY_121"
Reduce Output Operator
Size: 8
Expected Rows: 1
columnExprMap: {"VALUE._col0", "_col0"}
null sort order: ""
sort order: ""
value expressions: "_col0 (type: bigint)"
OperatorId: "RS_122"
outputname: "Reducer 3"
outputOperator: ["GBY_123"]
CUSTOM_SIMPLE_EDGE
Reducer 3
(vectorized=true)
Group By Operator
Size: 8
Expected Rows: 1
aggregations: ["count(VALUE._col0)"]
mode: "mergepartial"
outputColumnNames: ["_col0"]
OperatorId: "GBY_123"
File Output Operator
Size: 8
Expected Rows: 1
compressed: "false"
table: "{"input format":"org.apache.hadoop.mapred.SequenceFileInputFormat", "output format":"org.apache.hadoop.hive.ql.io.HiveSequenceFileOutputFormat", "serde":"org.apache.hadoop.hive.serde2.lazy.LazySimpleSerDe"}"
OperatorId: "FS_124"