

Build your own Database

Week 9

Assorted Findings

```
auto op = TableScanOp<T>(_operator, _search_value, _search_value_2);  
// [...]  
if (op(boost::get<T>((*column)[offset]))) {  
    _pos_list->push_back(pos);  
}
```

```
void handle_reference_column(ReferenceColumn& column,  
std::shared_ptr<  
    ColumnVisitableContext> context) {  
    // [...]  
    for (size_t virtual_id = 0; virtual_id < r_column.size();  
        ++virtual_id) {  
        if (_visit_operator(type_cast<T>(r_column[virtual_id]), _value,  
            _value2)) {  
            casted_context->appendPosition((*pos_list)[virtual_id]);  
        }  
    }  
}
```

Assorted Findings

```
void TableScan::execute() {
    const size_t column_id = _input_left->column_id_by_name(_filter_column_name);

    const std::string type_string = _input_left->column_type(column_id);
    auto column_instance = opossum::ValueColumn<type>(_input_left, column_id, type_string);
    const auto& instance = column_instance;
    const auto& type_id = instance.type_id();
    const auto& type_hash = instance.type_hash();

    if (type_hash == typeid(int32_t).hash()) {
        _out = TableScanImpl::execute(_input_left, column_id, type_string);
    } else if (type_hash == typeid(int64_t).hash()) {
        _out = TableScanImpl::execute(_input_left, column_id, type_string);
    } else if {
        // [...]
    } else {
        throw std::runtime_error("Unimplemented column_type");
    }
}

void TableScan::execute() {
    _input_left = _in->get_output();
    _filter_column_id = _input_left->column_id_by_name(
        _filter_column_name);
    auto column_type = _input_left->column_type(_filter_column_id);

    if (column_type == "int") {
        typed_execute<int32_t>();
    } else if (column_type == "long") {
        typed_execute<int64_t>();
    } else if (column_type == "float") {
        typed_execute<float>();
    } else if (column_type == "double") {
        typed_execute<double>();
    } else if (column_type == "string") {
        typed_execute<std::string>();
    } else {
        Fail("Unimplemented column_type");
    }
}
```

```

1  * This file contains the actual filter logic.
2  * Every filter has its own struct.
3  * The structs implement three major methods:
4
5  * check_value
6  *   This method is used to compare plain values (i.e. in
7  *   check_value_id
8  *   This method is used to compare value ids (i.e. in D)
9  *   Note that the comparison operator in use might be di
10 *   This will be explained in detail later.
11 * begin_dictionary_column
12 *   Since tables may have multiple chunks, and dictionaries
13 *   on a per-chunk basis, the value id of the filter val
14 *   This method is used to look up the respective value
15
16
17 * Optimizations
18 *
19 * Sorted, dictionary-compressed columns offer a great way
20 * First, we use binary searches to look up the respective
21 * Second, depending on the operator, we either use a lower
22 * The idea is to make use of the respective characteristics
23 * lower_bound
24 *   Returns the first value in a vector that is greater
25 *   Returns vector.end() if last value is strictly less
26 * upper_bound
27 *   Returns the first value in a vector that is greater
28 *   Returns vector.end() if last value is less than or
29 *
30 * In conclusion, this offers the following possibilities:
31 * Operator | Applied Logic
32 * -----
33 * == | lo / ==
34 * > | sb / >
35 * < | lb / <
36 * <= | sb / <
37 *
38 * As an example, let's look at the '>' operator.
39 * We use upper_bound to search for the value in the dict.
40 * We now have two options:
41 * 1. The searched value is in the dict.
42 *   upper_bound will return the value in the vector that
43 *   We can therefore include this value when we filter.
44 *   However, we do not include the searched value as th
45 * 2. The searched value is not in the dict.
46 *   upper_bound will return the value in the vector that
47 *   that is smaller than the searched value.
48 *   This value must be greater than the searched value.
49 * Consequently, using the '==' operator on the found value
50 *
51 * The main advantage we get out of this is that if the val
52 * we do not have to spend time to decide that we actually
53 * rather than the requested '>' operator.
54 * The other operators mentioned above behave similarly.
55 * The 'BETWEEN' operator is a combination of '>' and '<='
56 * '==' and '!=' use lower_bound and check if the returned v
57 *
58 * Additionally, the operators implement logic to recognize
59 * For example, if there is an equal scan requested on a di
60 * present in the dictionary, we can completely disregard t
61 *
62
63 #pragma once
64
65 #include <limits>
66 #include <vector>
67
68 #include "types.hpp"
69
70 namespace operator {
71
72 enum class ScanType { ALL, SCAN, NONE };
73
74 template <typename T>
75 struct EqFilter {
76     explicit EqFilter(const T &value) : value(value) {}

```

```

200 *
201 *
202 *
203 *
204 *
205 *
206 *
207 *
208 *
209 *
210 *
211 *
212 *
213 *
214 *
215 *
216 *
217 *
218 *
219 *
220 *
221 *
222 *
223 *
224 *
225 *
226 *
227 *
228 *
229 *
230 *
231 *
232 *
233 *
234 *
235 *
236 *
237 *
238 *
239 *
240 *
241 *
242 *
243 *
244 *
245 *
246 *
247 *
248 *
249 *
250 *
251 *
252 *
253 *
254 *
255 *
256 *
257 *
258 *
259 *
260 *
261 *
262 *
263 *

```

```

/ **
 * Get the right operation for the given operator.
 * For most operations, there is the possibility that either all values or no values match.
 * We can catch and easily process these cases by looking at the value that is returned by lower_bound or
 * upper_bound.
 * Say we have the following dictionary vectors:
 *
 * ValueID | Value
 * -----
 * 0 | B
 * 1 | C
 * 2 | D
 * 3 | F
 * 4 | G
 *
 * Then upper_bound (U) / lower_bound (L) return the following values:
 *
 * Value | U | L
 * -----
 * A | 0 | 0
 * B | 1 | 0
 * C | 3 | 2
 * D | 3 | 3
 * E | INV. | 4
 * F | INV. | INV.
 *
 * Then the table scan should return all values that match the following:
 * (X = No values, A = All Values)
 *
 * Operation | A | B | C | D | E | F | G | H |
 * -----
 * = | X | = 0 | = 2 | X | = 4 | X |
 * != | A | != 0 | != 2 | A | != 4 | A |
 * > | A | > 0 | > 2 | > 2 | X | X |
 * < | A | X | < 3 | < 3 | < 4 | X |
 * >= | A | A | >= 2 | >= 3 | >= 4 | X |
 * <= | X | = 2 | = 3 | = 3 | A | X |
 *
 * We then just pick the right method, according to the upper tables, and check for edge cases
 * (thus, an A or X in the table above). Afterwards, we iterate over the attribute vector and execute
 * the regarding method on it.
 */
std::pair<std::function<bool(ValueID)>, Match> get_dictionary_comparator(
    const std::string &op, const AllTypeVariant &allTypeVariant, const optional<AllTypeVariant> &allTypeVariant2,
    const DictionaryColumnsT &columns) const {
    const T value = type_cast<T>(allTypeVariant);

    // Calculate operation to check for valid entries.
    if (op == "=") {
        auto valueID = column.lower_bound(value);
        if (valueID != INVALID_VALUE_ID && column.value_by_value_id(valueID) == value) {
            return std::make_pair(valueID)(ValueID entry) { return entry == valueID; }, Match::one;
        } else {
            // In case we found did not find a value id that matches the given value,
            // we can assume that no entries with this value exist -> return an empty position list.
            return std::make_pair_none_match, Match::none;
        }
    } else if (op == "!=") {

```

Assorted Findings

```
template <typename T>
class ValueColumn : public BaseColumn {
    // [...]
    template <typename Operator>
    PosList typed_table_scan(const T& filter_value, const ChunkID
                            chunk_id, const PosList* in_pos_list =
                            nullptr) const;

    PosList table_scan_between(const T& filter_value, const T&
                               filter_value_2, const ChunkID chunk_id,
                               const PosList* in_pos_list) const;
};
```

Assorted Findings

```
void Projection::execute() {
    // [...]
    for (size_t chunk = 0; chunk < _table_in->chunk_count(); ++chunk) {
        // [...]
        for (size_t col = 0; col < _table_in->col_count(); ++col) {
            if (std::count(_columns.begin(), _columns.end(), _table_in->
                column_name(col)) > 0) {
                // [...]
                _table_out->get_chunk(chunk).add_column(_table_in->get_chunk(
                    chunk).get_column(col));
            }
        }
    }
}
```

Zeitplan

Woche	Phase	Mittwoch	Donnerstag
1	Sprint 1	19.10.2016	20.10.2016
2		26.10.2016	27.10.2016
3	Sprint 2	2.11.2016	3.11.2016
4		9.11.2016	10.11.2016
5	Sprint 3	16.11.2016	17.11.2016
6		23.11.2016	24.11.2016
7	Gruppenphase	30.11.2016	1.12.2016
8		7.12.2016	8.12.2016
9		14.12.2016 6	15.12.2016
10		4.1.2017	5.1.2017
11		11.1.2017	12.1.2017
12		18.1.2017	19.1.2017
13		25.1.2017	26.1.2017
14		1.2.2017	2.2.2017
15		8.2.2017	9.2.2017