

Develop your own Database

Week 8

Agenda

- Review Sprint 3
- Group Meetings
 - Remember to find a replacement for next week's meeting

Review

```
1 // Is used to do the actual comparsion of values.
2 auto matches_scan_criterion =
3     [](const auto& existing_value, const auto& search_value,
4         const ScanType& scan_type) -> bool {
5         switch (scan_type) {
6             case ScanType::OpEquals:
7                 return existing_value == search_value;
8             [...]
9         }
10    }
```

Review

```
1 const auto chunk = referenced_table->get_chunk(row_id.chunk_id);
2 const auto segment = chunk.get_segment(_column_id);
3 if ((value_segment = std::dynamic_pointer_cast<ValueSegment<T>>(segment)) {
4     values = value_segment->values();
5     // [...]
```

Review

```
1 ValueID find_in_dict(T value) const {
2     auto upper_bound_ref = std::find(_dictionary->cbegin(),
3                                     _dictionary->cend(),
4                                     value);
5
6     if (upper_bound_ref == _dictionary->cend()) {
7         return INVALID_VALUE_ID;
8     }
9     auto x = upper_bound_ref - _dictionary->cbegin();
10
11     return static_cast<ValueID>(x);
12 }
13 [...]
14 case ScanType::OpLessThan: {
15     auto value = typed_dict_segment->find_in_dict(search_value());
16     for (auto index = ValueID(0); index < size; index++) {
17         auto attribute = typed_dict_segment->attribute_vector()->get(i);
18         if (attribute < value) {
19             posList.emplace_back(RowID({i, ChunkOffset(index)}));
20         }
```

Review

```
1 for (size_t pos_list_id = 0; pos_list_id < size_of_pos_list; ++pos_list_id) {
2     RowID row_id = _pos_list->at(pos_list_id);
3
4     AllTypeVariant value = _referenced_table->get_value(row_id,
5                                                         _referenced_column_id);
6
7     resolve_data_type(data_type, [&](auto type) {
8         using Type = typename decltype(type)::type;
9         in_scope = scan_compare(scan_type,
10                                type_cast<Type>(value),
11                                type_cast<Type>(search_value));
```

Review

```
1 if (_add_all_chunk_offsets) {
2   // If we need to add all values of the source pos_list, we can just reference
the same list.
3   result_segment = std::make_shared<ReferenceSegment>(reference_source_segment-
>referenced_table(),
4                                                         reference_source_segment-
>referenced_column_id(),
5                                                         reference_source_segment-
>pos_list());
6 } else {
7   // We need to filter out some or all values. But, maybe for another segment
that shared the
8   // same PosList we already computed the resulting PosList for the new
Segment.
```

Review

```
1 for (ChunkID i = ChunkID(0); i < in_table->chunk_count(); i++) {
2 // [...]
3 case ScanType::OpEquals: {
4     auto value = typed_dict_segment->find_in_dict(search_value());
5
6     for (auto index = ValueID(0); index < size; index++) {
7         auto attribute = typed_dict_segment->attribute_vector()->get(i);
```


Review

```
[~/tmp/DYOD_WS1920] sprint3 ± ./scripts/lint.sh
```

std::function vs. lambda

- Seen many times:

```
1 std::function<bool(const T&)> _create_compare_function() {
2     switch (_scan_type) {
3         case ScanType::OpEquals:
4             return [=](const T& value) -> bool {
5                 return value == _search_value;
6             };
7         // [...]
```

std::function vs. lambda

std::function

Defined in header `<functional>`

```
template< class >  
class function; /* undefined */ (since C++11)
```

```
template< class R, class... Args >  
class function<R(Args...)>; (since C++11)
```

Class template `std::function` is a general-purpose polymorphic function wrapper. Instances of `std::function` can store, copy, and invoke any *Callable target* -- functions, lambda expressions, bind expressions, or other function objects, as well as pointers to member functions and pointers to data members.

- <http://quick-bench.com/AP81foBSotuVef9vIkDRYwvCn9Y>

std::function vs. lambda

„How to create lambdas depending on the comparator?“

```
1 enum class Condition {
2     Equals,
3     NotEquals
4 };
5
6 auto get_comparator(Condition condition) {
7     switch(condition) {
8         case Condition::Equals:
9             return [](auto value) {
10                 return value == 17;
11             };
12         case Condition::NotEquals:
13             return [](auto value) {
14                 return value != 17;
15             };
16     }
17 }
```

comp.cpp:13:7: error: 'auto' in return type deduced as '(lambda at comp.cpp:13:14)' here but deduced as '(lambda at comp.cpp:9:14)' in earlier return statement
return [](auto value) {

std::function vs. lambda

„How to create lambdas depending on the comparator?“

- You don't. Instead, pass whatever you want to do to the lambda
- Hollywood Principle / Inversion of Control

std::function vs. lambda

```
1  template <typename Functor>
2  void compare(Condition condition, Functor functor) {
3      switch(condition) {
4          case Condition::Equals:
5              functor(std::equal_to<void>{});
6              break;
7          case Condition::NotEquals:
8              functor(std::not_equal_to<void>{});
9              break;
10     }
11 }
12
13 void scan(const std::vector<int>& vector) {
14     const auto search_value = 17;
15     compare(Condition::Equals, [&](const auto comparator) {
16         for (const auto value : vector) {
17             if (comparator(value, search_value)) {
18                 std::cout << "match: " << value << std::endl;
19             }
20         }
21     });
22 }
```

This avoids the performance issues of std::function.

It was not expected for Sprint 3.

std::function vs. lambda

<http://quick-bench.com/tqljWKibQwNium5iZmmjLAFmVMQ>

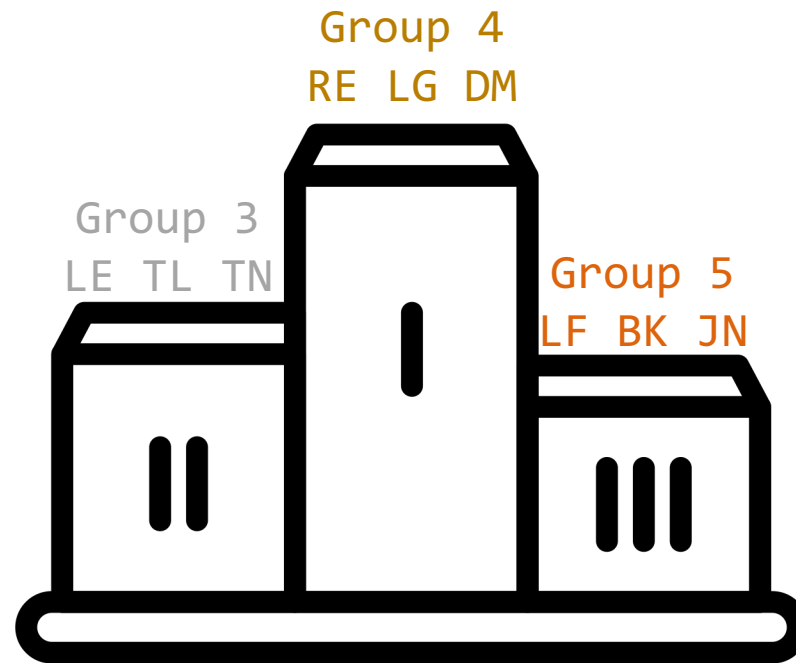
std::function vs. lambda

- Virtual method calls do not come for free
- In hot loops, we want to avoid them
- Trade-Off between performance and architecture

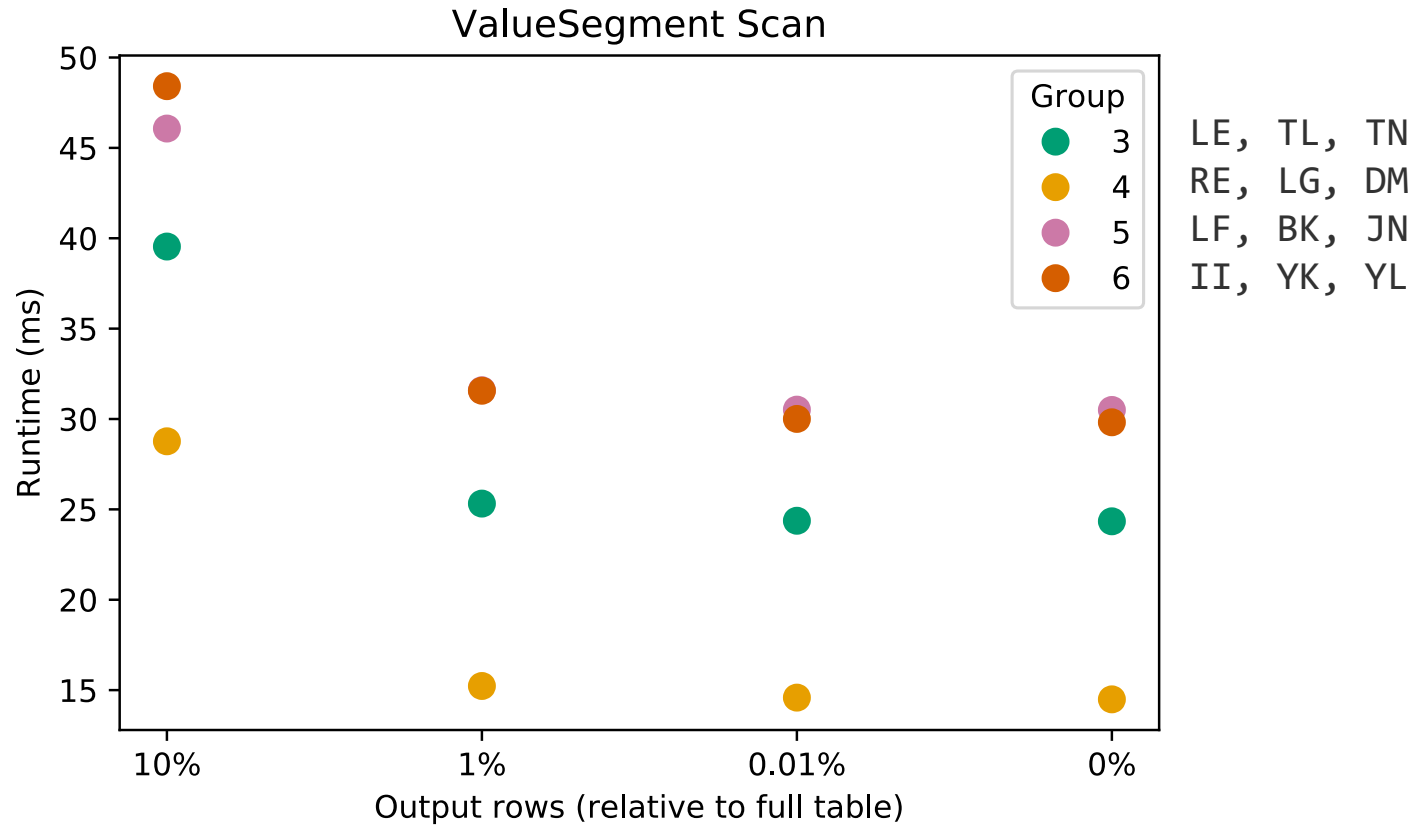
Performance Challenge

- 9 micro benchmarks to determine the performance of your table scans
 - Table size: 10M entries
 - Chunk sizes: **100K**, 1M, 10M
 - On Value-, Dictionary-, and ReferenceSegments
 - Predicate selectivities of 10%, 1%, 0.01%, 0%
 - We cleared the CPU cache between the experiments


Performance Challenge



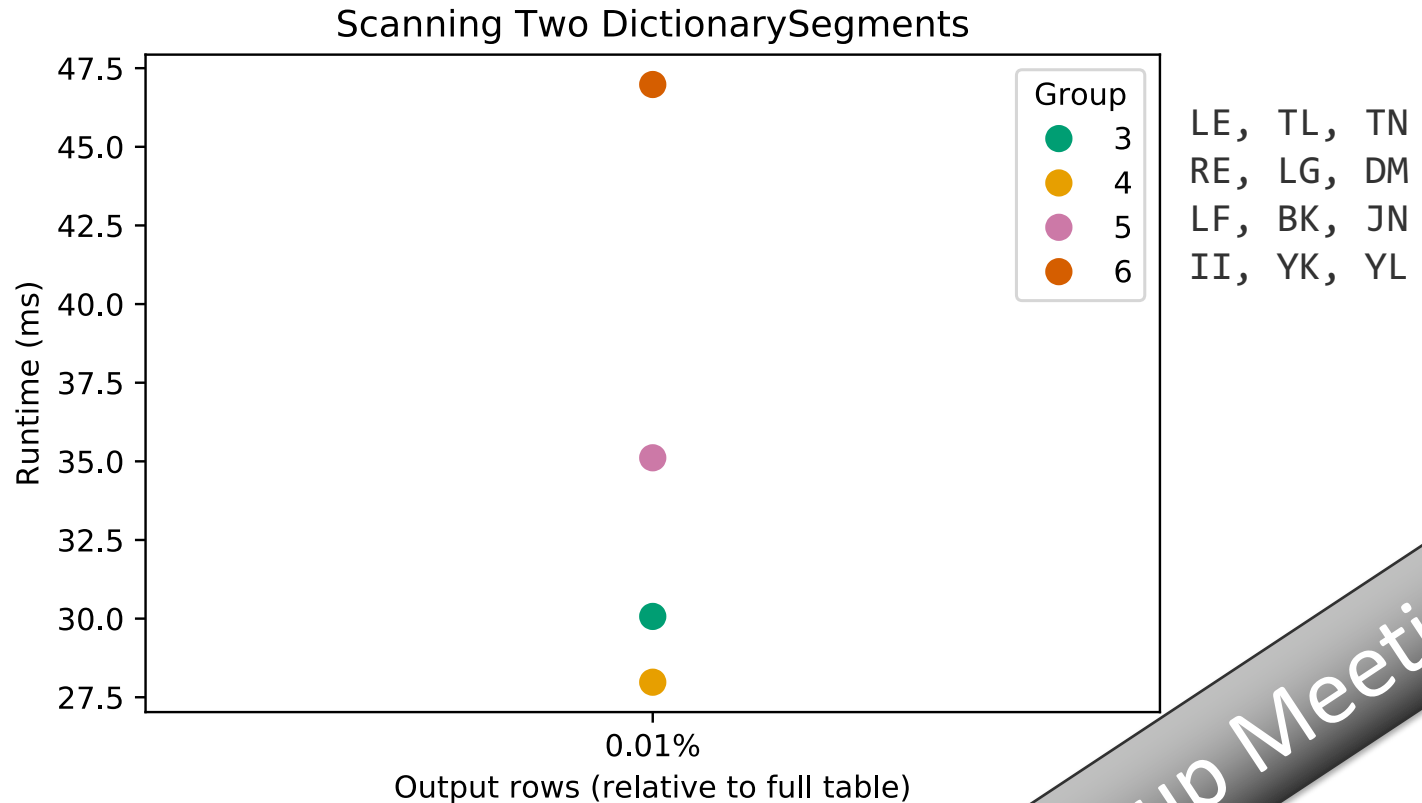
Performance Challenge



Performance Challenge



Performance Challenge



Next: Group Meetings