

**Final Presentation** 

Jekyeom Jeon Taoxi Jiang Yuttapichai (Guide) Kerdcharoen

### Re-iterate the motivation...



Source: <u>https://www.youtube.com/watch?v=rhxd\_xaeMPU</u>

# What we did (in a slide)

CMU-DB's standalone 128-bit fixed-point decimal	<ul> <li>Documented libfixeypointy</li> <li>Hardened libfixeypointy</li> <li>Improved libfixeypointy's multiplication and division performance</li> <li>Evaluated libfixeypointy against other standalone decimal implementations</li> </ul>	



#### All the code are documented



### More libfixeypointy boundary cases are handled



Decimal::Divide(const Decimal &denominator, const ScaleType &scale 1. Multiply the divid Division (by Tero scale), with overflo 2. If overflow, divide by the denominator with multi-word 256-bit 3. If no overflow, divide by the denominator with magic numbers in Moreover, the result is in the numerator's scale for technical rea If the result were to be in the denominator's scale, the first sta 10^(2\*denominator scale - numerator scale) which requires 256-bit

```
If value is 0
(denominator.ToNative() == 0) {
    rrow std::runtime_error("Divided by 0");
```

```
Decimal::Add(const Decima & other) {
    If both values are positive, if is possible to get overflowed
    (value_ > 0 && other.value_ > 0) {
    // Compute the maximum and safe value for other
    // E.g., 63 - 62 = 1
    // So, if other = 1, 62 + 1 = 63 (safe)
    // But if other = 2 (> 1), 62 + 2 = 64 (overflowed)
    int128_t other_bound = std::numeric_limits<__int128>::max() - valu
    if (other.value_ > other_bound) {
        throw std::runtime_error("Result overflow > 128 bits");
    }
```

# The performance bottleneck is grade-school multiply

#### Observation

- Both multiplication and division (by magic number) use 128-bit grade-school multiplication
- An existing grade-school multiply implementation contains a number of loops and potential bubbles

	128-bit Grade-School Multiply		128-bit Grade-School Mult	libfi libfi	
	libfixeypointy::Decimal::CalculateMultiWordProduct128	lib libfixeyp	inty::Decimal::DivideByConstantPowerOfTen128	lib	
I., lib., libfixeypointy	:Decimal::MultiplyAndSet				
libfixeypointy::Decimal::Multiply					
Bench					

Assumption
------------

• Unwinding the loop and manually reordering instructions (to avoid bubbles) could improve multiplication and division performance

# Optimizing by unwinding loops and reordering instructions

1. Let m and n = 2 and unwind loops		2. Reorder/remove instructions		
<pre>uint128_t k, t; uint32_t i, j; constexpr const uint128_t bo // Initialize first m chunks for (i = 0; i &lt; m; i++) half // For each chunk in b for (j = 0; j &lt; n; j++) { k = 0; // Match with all chunks i for (i = 0; i &lt; m; i++) { // Product + Old Value ( t = half_words_a[i] * ha // Take only bottom 64 b</pre>	to wi wo .f_ .ts	<pre>uint128_t k, t; constexpr const uint128_t bottom_mask = // Initialize first m chunks with 0 half_words_result[0] = 0; half_words_result[1] = 0; // for each chunk in b // j = 0, i = 0 t = half_words_a[0] * half_words_b[0] + half_words_result[0] = t &amp; bottom_mask; k = t &gt;&gt; 64; // j = 0, i = 1 t = half_words_result[1] = t &amp; bottom_mask; k = t &gt;&gt; 64; // End half_words_result[2] = k; // j = 1 k = 0; // j = 1 k = 0;</pre>	(uint nalf_	<pre>void _CalculateMultiWordProduct128_2_2(const uint128_t *const half_words</pre>
<pre>half_words_result[i + j]     // Carry     k = t &gt;&gt; 64; } half_words_result[j + m] =</pre>	= k;	<pre>// j = 1, 1 = 0 t = half_words_a[0] * half_words_b[1] + half_words_result[1] = t &amp; bottom_mask; k = t &gt;&gt; 64; // j = 2, i = 1 t = half_words_a[1] * half_words_b[1] + half_words_result[2] = t &amp; bottom_mask; k = t &gt;&gt; 64; // f = 4</pre>	nalf_ nalf_	<pre>t3 += (t2 &gt;&gt; 64); half_words_result[3] = (t3 &gt;&gt; 64); half_words_result[1] = t2 &amp; bottom_mask; half_words_result[2] = t3 &amp; bottom_mask;</pre>
}		half words result[3] = k;		}

#### Grade-school multiply is no longer bottleneck



# Division

- Specify custom predefined magic number to speed up division (and multiplication)
- Not too many of them (depending on the operations you normally want to do)
- No magic number -> predicate small -> hot path taking the branch (good)
- A few magic number -> hot path not taking the branch -> predicate small (good)
- Lots of magic number -> predicate big -> access pattern uniform anyway, doesn't make sense to add those magic number (bad)
- Generate and cache all seen magic numbers? -> will test in the future

Todo: Macro to convert static hashtable lookup to compiled predicates

#### ... div by zero / power of 2 check

- / 2. If not possible, regular division.
- if (MAGIC\_CUSTOM\_128BIT\_CONSTANT\_DIVISION.count(constant) == 0) {
   value\_ = static\_cast<uint128\_t>(value\_) / constant;
   return;

#### ... Magic Number Division



# Verification

- Python decimal (hybrid of fixed point and float)
  - Random **op1 +-\*/ op2**, repeated millions of times
  - Compare rounded off values
- Java BigDecimal
  - Long random chain **op1 +-\*/ op2 +-\*/ op3 …**
  - Compare error handling behavior (overflow), report and revert to previous value when error encountered
  - Compare exact values
  - Results exactly match



Operations on decimals stored in 32bit size, most digits **<u>before</u>** decimal point (scale is small) Dataset fits in L3 cache. Larger dataset results will come in future.





128bit Op Time (10000 line, 1000 iteration, avg 5 runs, skip first 2)



128bit Op Time (10000 line, 1000 iteration, avg 5 runs, skip first 2)



# Integrating libfixeypointy as PostgreSQL's UDT





Op Time (4M tuples, 2 columns, 5 runs avg)

Libfixeypointy(24B) NUMERIC(Variable Size) DOUBLE(8B) REAL(4B)



Op Time (4M tuples, 2 columns, 5 runs avg)

Libfixeypointy(24B) NUMERIC(Variable Size) DOUBLE(8B) REAL(4B)



pgfixeypointy is slower than NUMERIC (ToString() is slow)

Op Time (4M tuples, 2 columns, 5 runs avg)

Libfixeypointy(24B) NUMERIC(Variable Size) DOUBLE(8B) REAL(4B)



Op Time (4M tuples, 2 columns, 5 runs avg)

Libfixeypointy(24B) NUMERIC(Variable Size) DOUBLE(8B) REAL(4B)



## **Future Work**

- Support variable size: store a small decimal in a 64-bit or 32-bit
- Parallel aggregation to improve throughput
- > Running perf to find an opportunity for optimizing multiplication performance
- > Improve result writing (probably, string conversion) performance
- ➢ More Aggregator support: AVG, STD, VAR
- > Type Casting: Operations between different types (e.g., double+libfixeypointy)
- > More realistic workloads

## Resources

- libfixeypointy <u>https://github.com/cmu-db/libfixeypointy/tree/develop</u>
- pgfixeypointy <u>https://github.com/pnxguide/pgfixeypointy</u>



Prompt: a fast stream of 128-bit fixed-point decimals