

# Exploring Query Execution Strategies for JIT, Vectorization and SIMD

Tim Gubner, Peter Boncz

CWI

September 1, 2017

# Motivation

- What will future DB engines look like?
- New hardware increasingly heterogenous: GPUs, FPGAs on the same chip, specialized CPU instructions
- For performance: exploit hardware features
- → Multi-dimensional design space

# Contribution

- Shed some light into the design space
- Compare multiple implementations of TPC-H Q1:
  - ▶ Tuple-at-a-time / block-at-a-time
  - ▶ Regular / compact data types
  - ▶ Overflow checking / prevention
  - ▶ Different aggregation techniques
  - ▶ Varying aggregation table layout

# SIMD (Single Instruction Multiple Data)

- One instruction processes 512 bits (AVX-512)
- Less bits per element = more elements processed per instruction

<b>Data type width</b>	<b>Parallelism</b>
64-bit	8
32-bit	16
16-bit	32
8-bit	64
1-bit	512

# Data types

- Derived from schema:  
DECIMAL(15,2)  $\rightarrow$  64-bit integer ( $32 < \log_2(10^{15}) < 64$ )
- Thinner data types  $\rightarrow$  more data items per clock cycle

## Compact data types by example

- Reduce size of data types based on actual data
- Example from TPC-H:
  - ▶ Schema:  $tax \in DECIMAL(15, 2) \rightarrow$  64-bit integer
  - ▶ Data:  $tax \in \{0.0, 0.1, \dots, 0.8\} \rightarrow [0, 80] \rightarrow$  8-bit integer
  - ▶ Computation:  $tax \cdot tax \rightarrow$  8-bit integer  $\cdot$  8-bit integer = 16-bit integer
- Restrict data types using domain minimum & maximum

# Overflow handling

- Overflow handling is required to guarantee correctness

<b>Detection</b>	<b>Prevention</b>
Check via overflow CPU flag or specific code	Larger data types, overflow cannot (realistically) happen

- Overflow checking code is inefficient
- Overflow CPU flag is not SIMD friendly
- → Overflow prevention

## Case study: TPC-H Q1

```
SELECT
  l_returnflag, l_linestatus,
  count(*) AS count_order
  sum(l_quantity) AS sum_qty,
  avg(l_quantity) AS avg_qty,
  avg(l_discount) AS avg_disc,
  avg(l_extendedprice) AS avg_price,
  sum(l_extendedprice) AS sum_base_price,
  sum(l_extendedprice*(1-l_discount)) AS sum_disc_price,
  sum(l_extendedprice*(1-l_discount)*(1+l_tax))
  AS sum_charge
FROM
  lineitem
WHERE
  l_shipdate <= date '1998-12-01' - interval '90' day
GROUP BY
  l_returnflag, l_linestatus
ORDER BY
  l_returnflag, l_linestatus
```



## Compact data types in Q1

Expression	Regular	Compact
<code>l_tax</code>	64	8
<code>l_returnflag</code>	8	8
<code>l_linestatus</code>	8	8
<code>l_extendedprice</code>	64	32
<code>1+l_tax</code>	64	8
<code>(1-l_discount)*(1+l_tax)</code>	64	16
<code>l_extendedprice*(1-l_discount)</code>	64	32
<code>l_extendedprice*(1-l_discount)*(1+l_tax)</code>	64	64

## Q1 flavors


<b>Base flavor</b>	<b>Intermediate storage</b>	<b>Selection strategy</b>
X100 HyPer Handw. AVX-512	CPU cache Registers Registers	Sel. vector Branching Bit mask

# Design space explored

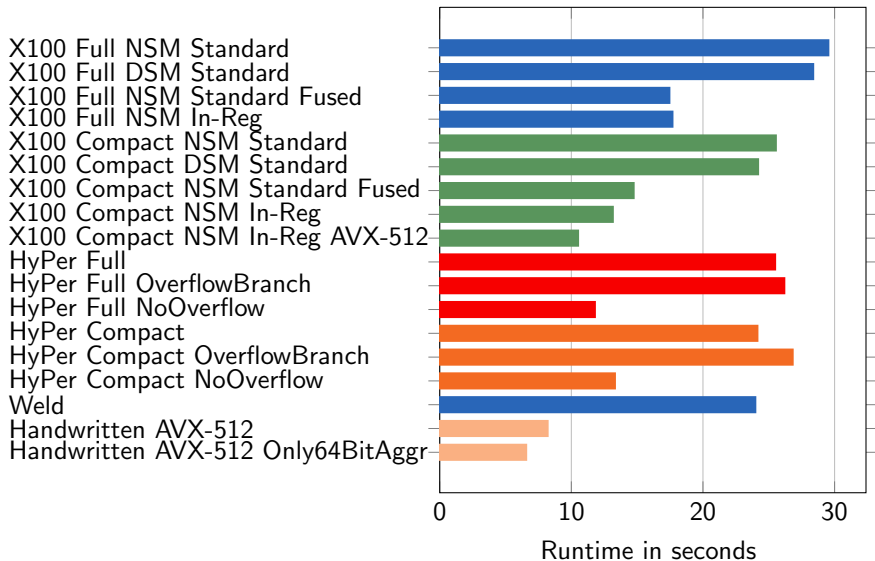
- X100
  - ▶ Data types: Full, Compact
  - ▶ Aggregate table layout: NSM (row-wise), DSM (column-wise)
  - ▶ Aggregation algorithm: Standard, Standard Fused, In-Reg
- HyPer
  - ▶ Data types: Full, Compact
  - ▶ Overflow: Default (set flag), OverflowBranch (branching), NoOverflow (prevention)

# Q1 flavors on Sandy Bridge

X100 Full NSM Standard  
X100 Full DSM Standard  
X100 Full NSM Standard Fused  
X100 Full NSM In-Reg  
X100 Compact NSM Standard  
X100 Compact DSM Standard  
X100 Compact NSM Standard Fused  
X100 Compact NSM In-Reg  
HyPer Full  
HyPer Full OverflowBranch  
HyPer Full NoOverflow  
HyPer Compact  
HyPer Compact OverflowBranch  
HyPer Compact NoOverflow  
Weld



# Q1 flavors on Knights Landing



# Takeaways

- Do not rely on schema!
- Exploit data statistics instead
- Ingredients for performance:
  - ▶ Compact data types
  - ▶ Overflow prevention
  - ▶ Adapt algorithms to take advantage of compact data types

## Future work

- Even thinner data types / different representations → compressed execution
- System that generates such a query

Questions?