


Designing an adaptive VM that combines vectorized and JIT execution on heterogeneous hardware

Tim Gubner

ICDE PhD Symposium, 2018



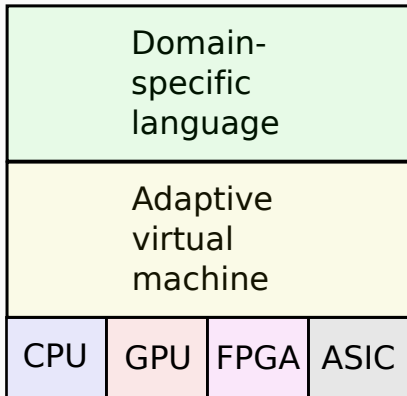
Centrum Wiskunde & Informatica



System	CPU	GPU	FGPA	ASIC
MonetDB	✓			
doppioDB	✓		✓	
Ocelot	✓	✓		
HyPer	✓			
MapD	✓	✓		
CoGaDB	✓	✓		
TensorFlow	✓	✓	?	✓

One system to rule them all

One system to bring them, and in Dark Silicon bind them



Virtual Machine

- Compilation is time consuming (≥ 20 ms ¹)
- Also noticeable in HyPer ²
- Compilers make assumptions.
Resulting code either:
 - Static and concise
 - Dynamic and bulky (code explosion)

¹Using LLVM C++ API and optimization passes

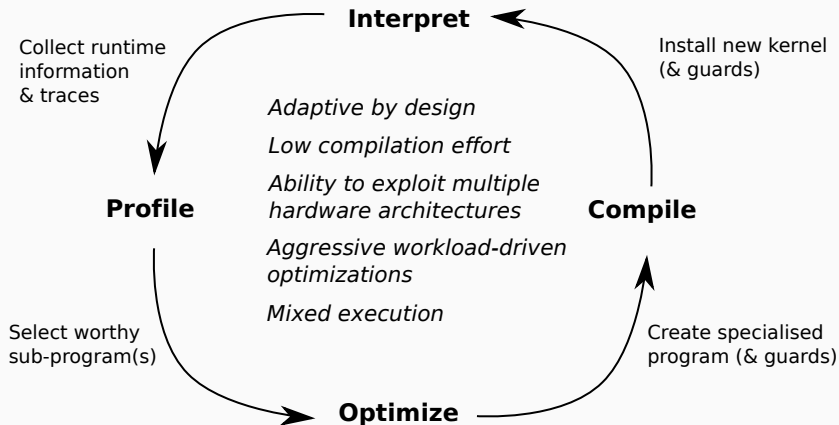
²Kohn et al. "Adaptive Execution of Compiled Queries", ICDE 2018

- Compilation is time consuming (≥ 20 ms ¹)
- Also noticeable in HyPer ²
- Compilers make assumptions.
Resulting code either:
 - Static and concise
 - Dynamic and bulky (code explosion)

Why would we ALWAYS want to compile EVERYTHING?

¹Using LLVM C++ API and optimization passes

²Kohn et al. "Adaptive Execution of Compiled Queries", ICDE 2018



Domain-Specific Language

Low enough

- Micro-adaptivity³
- Efficient interpretation
- JIT / incremental compilation

³Răducanu et al. "Micro adaptivity in Vectorwise", SIGMOD 2013

Low enough

- Micro-adaptivity³
- Efficient interpretation
- JIT / incremental compilation

High enough

- Efficient execution on multiple devices
- Macro-adaptivity: e.g. reorder operations

³Răducanu et al. "Micro adaptivity in Vectorwise", SIGMOD 2013

Low enough

- Micro-adaptivity³
- Efficient interpretation
- JIT / incremental compilation

High enough

- Efficient execution on multiple devices
- Macro-adaptivity: e.g. reorder operations

Goal

Relation algebra $\rightarrow ? \rightarrow$ Assembly, OpenCL ...

³Răducanu et al. "Micro adaptivity in Vectorwise", SIGMOD 2013

Relational algebra

Too high-level

(Scalar) Monad/Monoid comprehension

Weld ⁴, MRQL ⁵

High-level but per-tuple transformations lose information

⁴S. Palkar et al. "Weld: A Common Runtime for High Performance Data Analytics", CIDR 2017

⁵Fegaras, L. "An Algebra for Distributed Big Data Analytics", 2016

C alikes

OpenCL, CUDA, Intel SPMD (ispc) ...

Too low-level

MonetDB assembly language

Heavily data-parallel, too low-level

Data-parallelism as first-class citizen

- Data-parallel skeletons/patterns
Specialized operations on chunks of data
For example: `map`, `filter`, `scatter`, `gather` ...
- Lambda functions
- Immutable variables for intermediates (Static single assignment form)
- Mutable variables for remaining state
- Partially typed (`a ∈ DECIMAL(6, 2)` instead of `a ∈ int64_t`)

Op.	map	filter	scatter	gather	ht_ins	merge
π	✓					
σ	✓	✓				
\bowtie Hash	✓	✓	✓	✓	✓	
G Hash	✓	✓	✓	✓	✓	
\cup Hash	✓	✓	✓	✓	✓	
\bowtie Merge			✓	✓		✓
Sort			✓	(✓)		✓

Skeletons themselves do not need to be implemented data-parallel (e.g. `ht_ins`)...

```
mut i
mut k
i := 0
k := 0
loop
  let input = read i some_data in
    let a = map (\x -> 2*x) input in
      let t = filter (\x -> x>0) a in
        let b = condense t
          write x i a
          write y k b
          i := i + len(a)
          k := k + len(b)

if i >= 4096 then
  break
```

Plan

Base framework

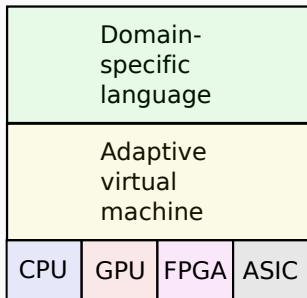
DSL, vectorized interpreter

Dynamic VM

Workload-specific optimizations

Multiple target architectures

GPUs, potentially FPGAs



DSL

- Abstract enough for:
 - Efficient portability
 - Adaptive optimizations
 - Efficient interpretation
- State of art does not fit!
- *Data parallelism as first-class citizen*

VM

- *Interpret first, maybe compile later*
- Cost-models are hard to get right!
- Adaptive by design
- Aggressive workload-driven optimizations
- Mixed execution