FRIEDRICH-ALEXANDER
UNIVERSITÄT
ERLANGEN-NÜRNBERG
FACULTY OF ENGINEERING

Automated Instruction Stream Throughput Prediction for Intel and AMD Microarchitectures

J. Laukemann, J. Hammer, J. Hofmann, G. Hager, G. Wellein

Overview

1. **Analytic Performance Modeling**
   - Why?
   - Components
   - What we do in this work
2. **Model Construction**
   - Model assumptions: port model, full throughput,…
   - Microbenchmarking for instruction throughput (and latency)
   - Putting together a prediction
3. **OSACA: Automating the in-core model construction**
   - Overview
   - Structure and Output
4. **Schönauer Triad Benchmark Example**
5. **π Benchmark Example**

Performance Modeling For Loop Kernels

- **How fast can my kernel run *at best*?**
- **What are the relevant hardware bottlenecks?**
- **Apply simplified model of underlying hardware**
  - In-core execution
  - Data transfer
  - Putting execution and data transfer together

Roofline Model

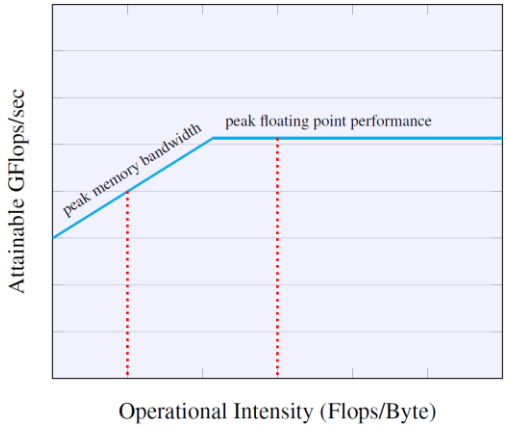

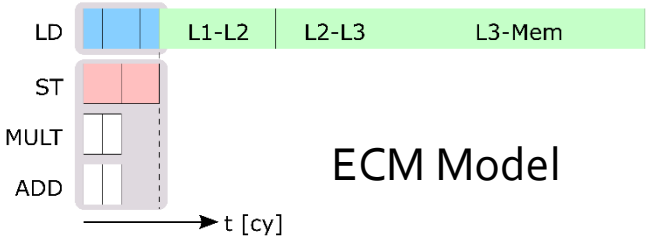

ECM Model

Benefits

- Optimization within the kernel
- Guiding decisions for or against specific architecture
- Deeper understanding of code and hardware interaction

This Work

- Semi-automated machine instruction (throughput/latency) benchmarking
- Automated in-core runtime prediction for steady-state loops
- Open-Source Architecture Code Analyzer (OSACA) tool
- Case studies

OSACA – Workflow

User input

```
// TRIAD BENCHMARK

//STARTLOOP
for(int j=0; j<size; ++j){
    a[j] = b[j] + c[j]*d[j];
}
```

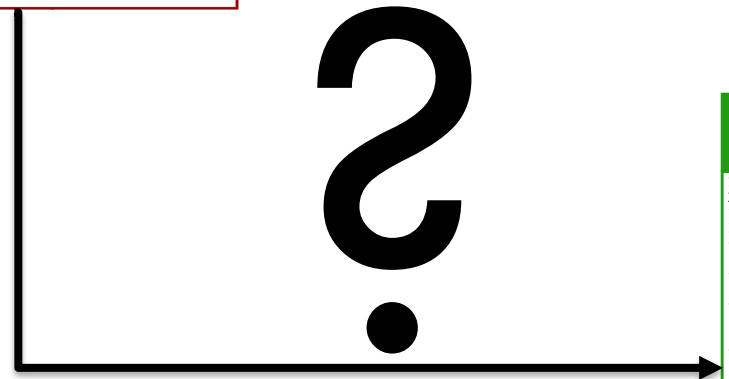

Throughput analysis

```
P - Load operation can be hidden behind a past or future store instruction
X - No information for this instruction in data file
* - Instruction micro-ops not bound to a port
Port Binding in Cycles Per Iteration:
---------------------------------------------------------------------------
| Port   |  0  |  1  |  2  |  3  |  4   |  5   |  6   |  7   |  8  |  9  |
---------------------------------------------------------------------------
| Cycles | 1.0 | 1.0 |  0  |  0  | 0.75 | 0.75 | 0.75 | 0.75 | 4.0 | 4.0 |
---------------------------------------------------------------------------

                        Ports Pressure in cycles
|  0   |  1   |  2   |  3   |  4   |  5   |  6   |  7   |  8   |  9   |
---------------------------------------------------------------------------
|      |      |      |      |      |      |      |      |      |      | X .L10:
|      |      |      |      |      |      |      |      | (1.0)| (1.0)| P vmovapd (%r15,%rax), %ymm0
|      |      |      |      |      |      |      |      | 1.00 | 1.00 | vmovapd (%r12,%rax), %ymm3
|      |      |      |      | 0.25 | 0.25 | 0.25 | 0.25 |      |      | addl $1, %ecx
| 1.00 | 1.00 |      |      |      |      |      |      | 1.00 | 1.00 | vfmadd132pd 0(%r13,%rax),
|      |      |      |      |      |      |      |      |      |      | %ymm3, %ymm0
|      |      |      |      |      |      |      |      | 2.00 | 2.00 | vmovapd %ymm0, (%r14,%rax)
|      |      |      |      | 0.25 | 0.25 | 0.25 | 0.25 |      |      | addq $32, %rax
|      |      |      |      | 0.25 | 0.25 | 0.25 | 0.25 |      |      | cmpl %ecx,%r10d
|      |      |      |      |      |      |      |      |      |      | ja .L10
Total number of estimated throughput: 4.0
```

OSACA – Workflow

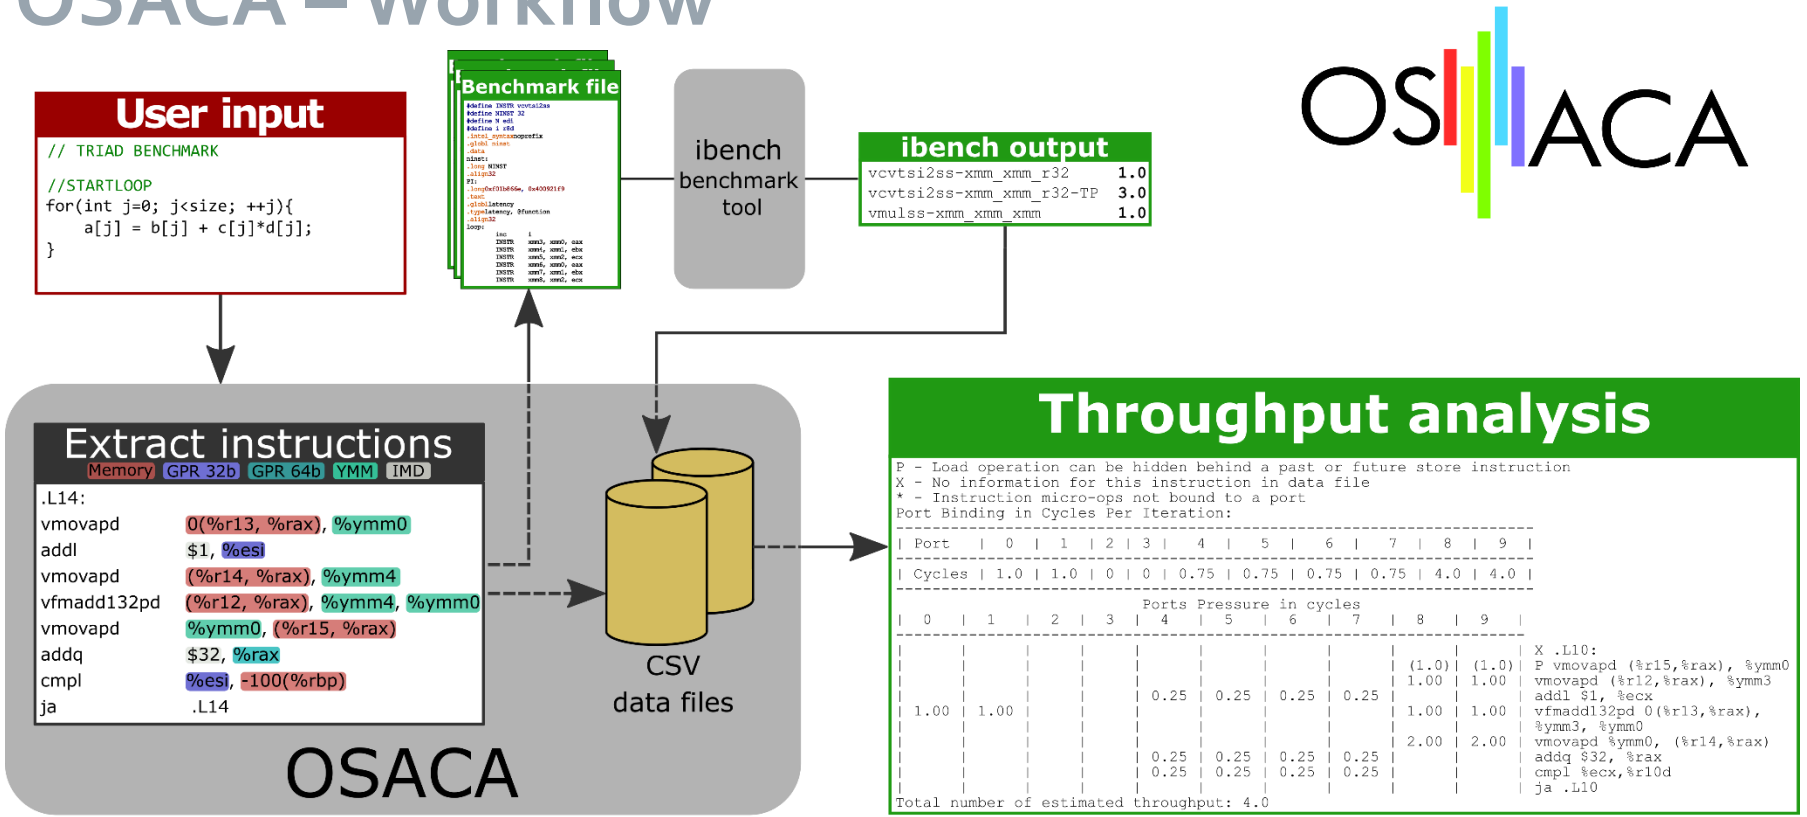

Model Construction (I): Assumptions

1. All Data in L1
2. Average distribution of port scheduling
3. Perfect out-of-order scheduling
4. Latencies hidden via speculative execution
5. Runtime prediction == longest time any port is occupied

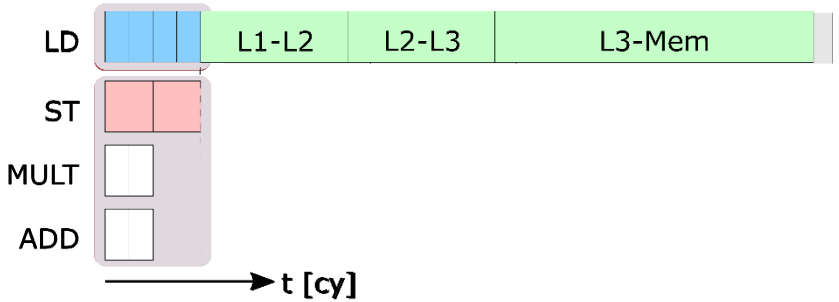

Model Construction (I): Assumptions

1. All Data in L1
2. Average distribution of port scheduling
3. Perfect out-of-order scheduling
4. Latencies hidden via speculative execution
5. Runtime prediction == longest time any port is occupied

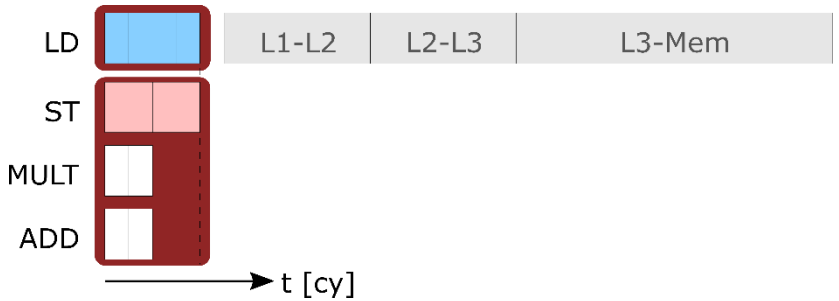

Model Construction (II): Port Models

Intel Skylake

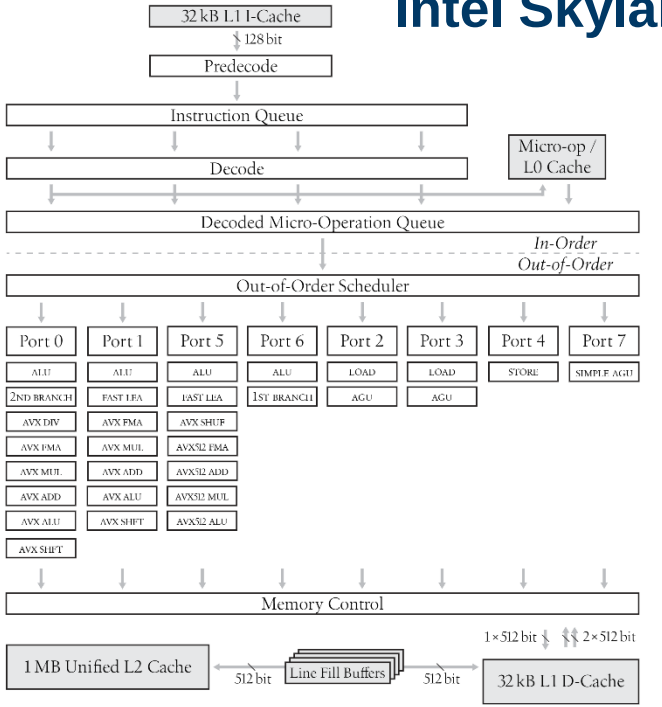

AMD Zen

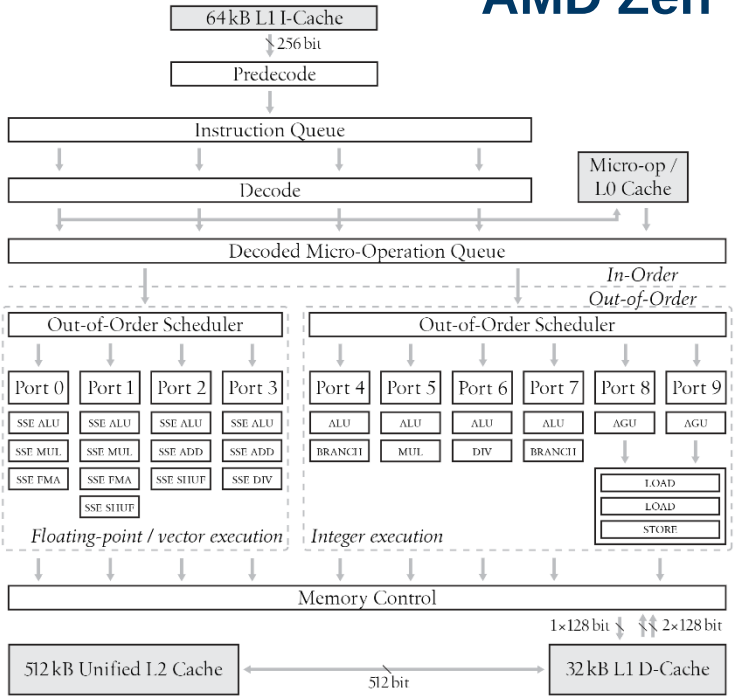

Model Construction (II): Port Models

Intel Skylake

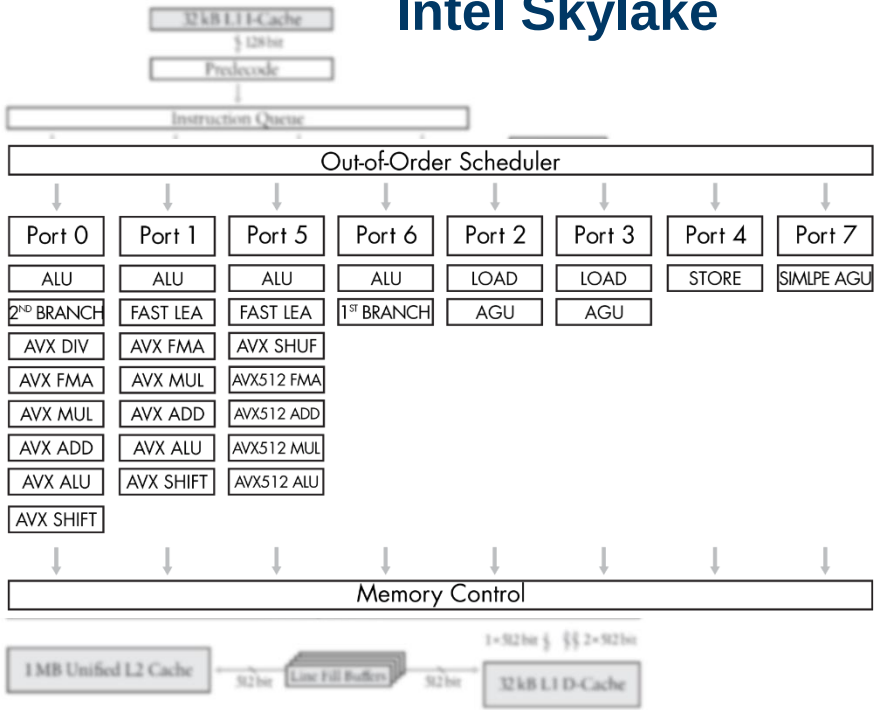

AMD Zen

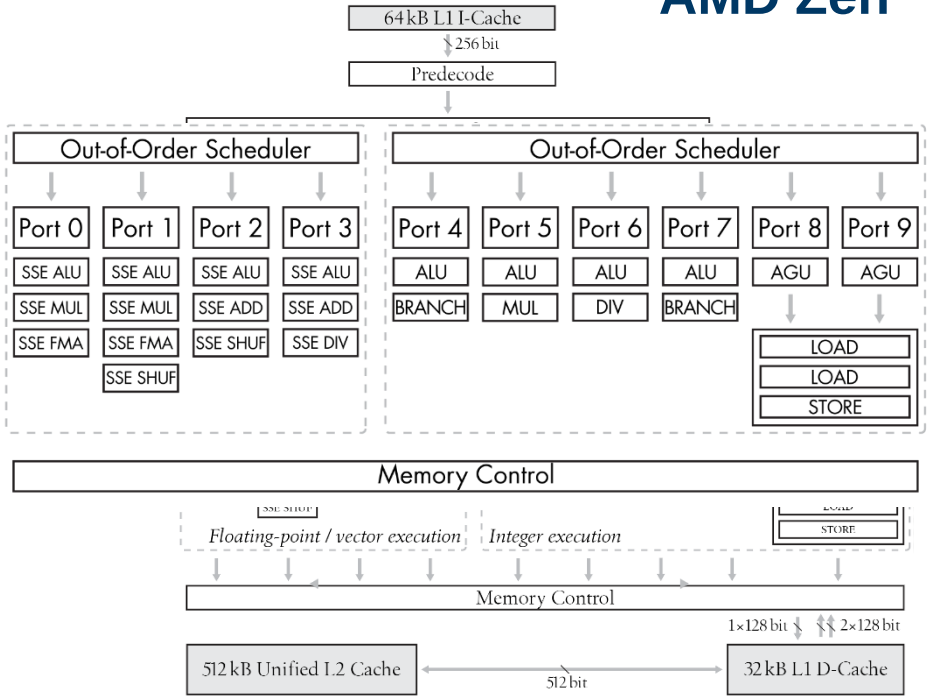

OSACA – Workflow

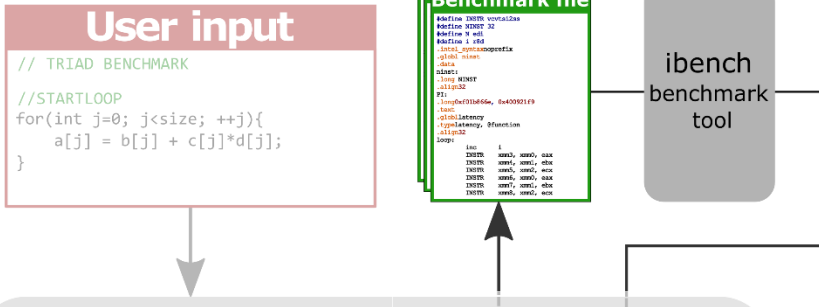

ibench benchmark tool

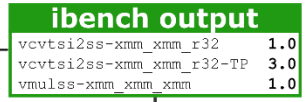

OSACA

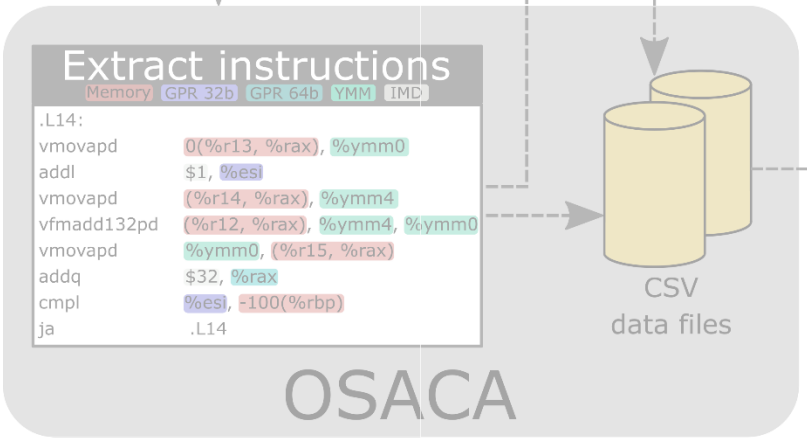

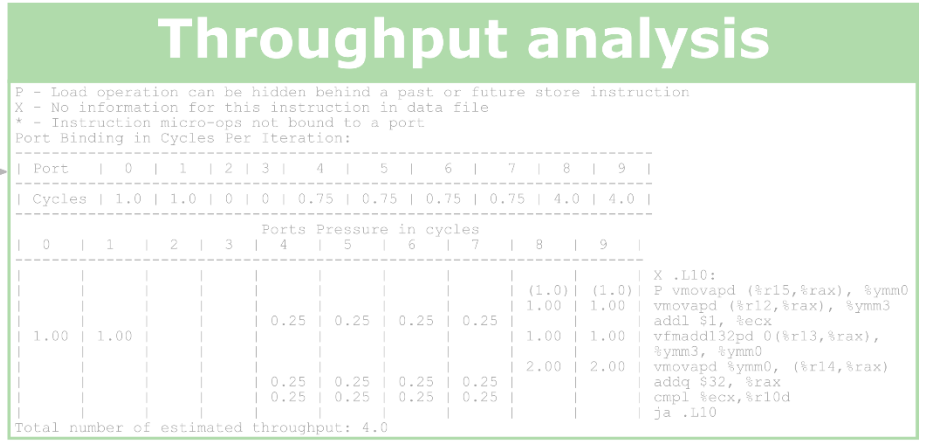

Model Construction (III): Microbenchmarks

Latency

```
loop:
  inc %eax
  vaddpd %xmm0, %xmm1, %xmm0
  vaddpd %xmm0, %xmm0, %xmm1
  vaddpd %xmm0, %xmm1, %xmm0
  ...
  vaddpd %xmm0, %xmm0, %xmm1
  cmp %eax, %edx     #loop count
  jl loop
```

Throughput

```
loop:
  inc %eax
  vaddpd %xmm0, %xmm0, %xmm3
  vaddpd %xmm1, %xmm1, %xmm4
  vaddpd %xmm2, %xmm2, %xmm5
  vaddpd %xmm0, %xmm0, %xmm6
  vaddpd %xmm1, %xmm1, %xmm7
  vaddpd %xmm2, %xmm2, %xmm8
  vaddpd %xmm0, %xmm0, %xmm9
  ...
  cmp %eax, %edx     #loop count
  jl loop
```

Model Construction (III): Microbenchmarks

Latency

```
loop:
  inc %eax
  vaddpd %xmm0, %xmm0, %xmm0
  vaddpd %xmm0, %xmm0, %xmm1
  vaddpd %xmm1, %xmm1, %xmm0
  ...
  vaddpd %xmm0, %xmm0, %xmm1
  cmp %eax, %edx     #loop count
  jl loop
```

FUTURE WORK

Throughput

```
loop:
  inc %eax
  vaddpd %xmm0, %xmm0, %xmm3
  vaddpd %xmm1, %xmm1, %xmm4
  vaddpd %xmm2, %xmm2, %xmm5
  vaddpd %xmm0, %xmm0, %xmm6
  vaddpd %xmm1, %xmm1, %xmm7
  vaddpd %xmm2, %xmm2, %xmm8
  vaddpd %xmm0, %xmm0, %xmm9
  ...
  cmp %eax, %edx     #loop count
  jl loop
```

Model Construction (III): continued

Benchmark tool output

```
Using frequency 1.80GHz.

vaddpd-xmm_xmm_xmm-1:          4.009
vaddpd-xmm_xmm_xmm-2:          2.006
vaddpd-xmm_xmm_xmm-4:          1.011
vaddpd-xmm_xmm_xmm-5:          0.805
vaddpd-xmm_xmm_xmm-8:          0.556
vaddpd-xmm_xmm_xmm-10:         0.554
vaddpd-xmm_xmm_xmm-12:         0.551
```

database entry

```
vaddpd-xmm_xmm_xmm, 0.5, 4.0, \
"(0.5,0,0.5,0,0,0,0,0,0)"
```

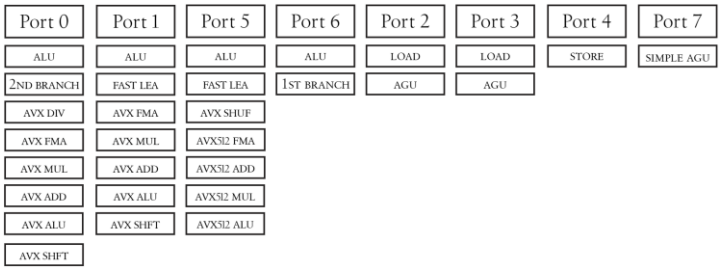

Model Construction (III): continued

Benchmark tool output

```
Using frequency 1.80GHz.

vaddpd-xmm_xmm_xmm-1:          4.009
vaddpd-xmm_xmm_xmm-2:          2.006
vaddpd-xmm_xmm_xmm-4:          1.011
vaddpd-xmm_xmm_xmm-5:          0.805
vaddpd-xmm_xmm_xmm-8:          0.556
vaddpd-xmm_xmm_xmm-10:         0.554
vaddpd-xmm_xmm_xmm-12:         0.551
```

# of independent instructions

database entry

```
vaddpd-xmm_xmm_xmm, 0.5, 4.0, \
"(0.5,0,0.5,0,0,0,0,0,0)"
```

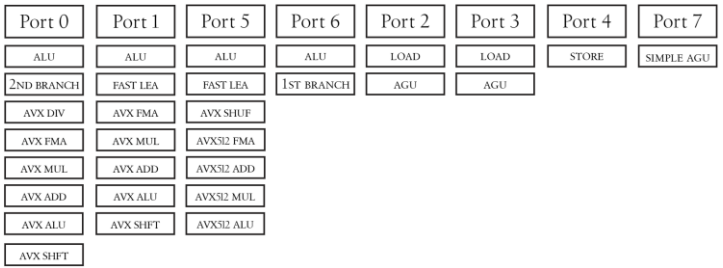

Model Construction (III): continued

Benchmark tool output

```
Using frequency 1.80GHz.

vaddpd-xmm_xmm_xmm-1:          4.009
vaddpd-xmm_xmm_xmm-2:          2.006
vaddpd-xmm_xmm_xmm-4:          1.011
vaddpd-xmm_xmm_xmm-5:          0.805
vaddpd-xmm_xmm_xmm-8:          0.556
vaddpd-xmm_xmm_xmm-10:         0.554
vaddpd-xmm_xmm_xmm-12:         0.551
```

# of independent instructions

CPI

database entry

```
vaddpd-xmm_xmm_xmm, 0.5, 4.0, \
"(0.5,0,0.5,0,0,0,0,0,0)"
```

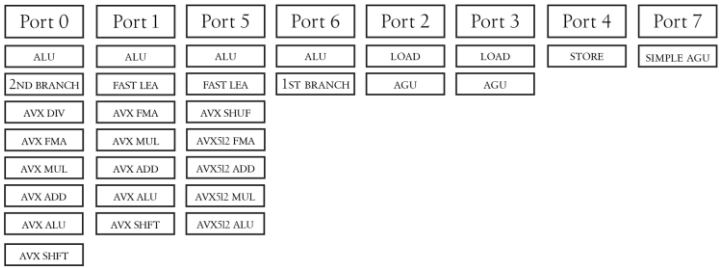

Model Construction (III): continued

Benchmark tool output

```
Using frequency 1.80GHz.

vaddpd-xmm_xmm_xmm-1:          4.009
vaddpd-xmm_xmm_xmm-2:          2.006
vaddpd-xmm_xmm_xmm-4:          1.011
vaddpd-xmm_xmm_xmm-5:          0.805
vaddpd-xmm_xmm_xmm-8:          0.556
vaddpd-xmm_xmm_xmm-10:         0.554
vaddpd-xmm_xmm_xmm-12:         0.551
```

# of independent instructions

CPI

database entry

```
vaddpd-xmm_xmm_xmm, 0.5, 4.0, \
"(0.5,0,0.5,0,0,0,0,0,0,0)"
```

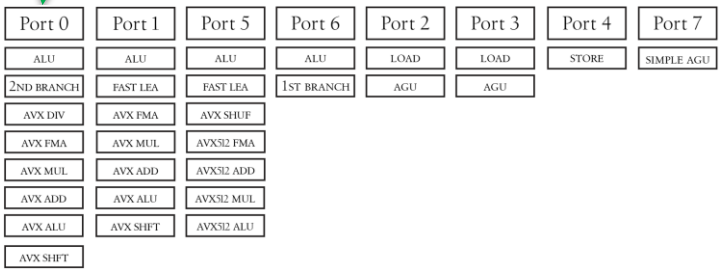

Model Construction (III): continued

Benchmark tool output

```
Using frequency 1.80GHz.

vaddpd-xmm_xmm_xmm-1:          4.009
vaddpd-xmm_xmm_xmm-2:          2.006
vaddpd-xmm_xmm_xmm-4:          1.011
vaddpd-xmm_xmm_xmm-5:          0.805
vaddpd-xmm_xmm_xmm-8:          0.556
vaddpd-xmm_xmm_xmm-10:         0.554
vaddpd-xmm_xmm_xmm-12:         0.551
```

# of independent instructions

CPI

database entry

```
vaddpd-xmm_xmm_xmm, 0.5, 4.0, \
"(0.5,0,0.5,0,0,0,0,0,0,0)"
```

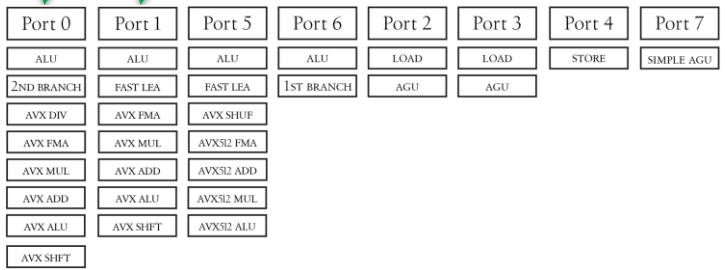

OSACA – Workflow

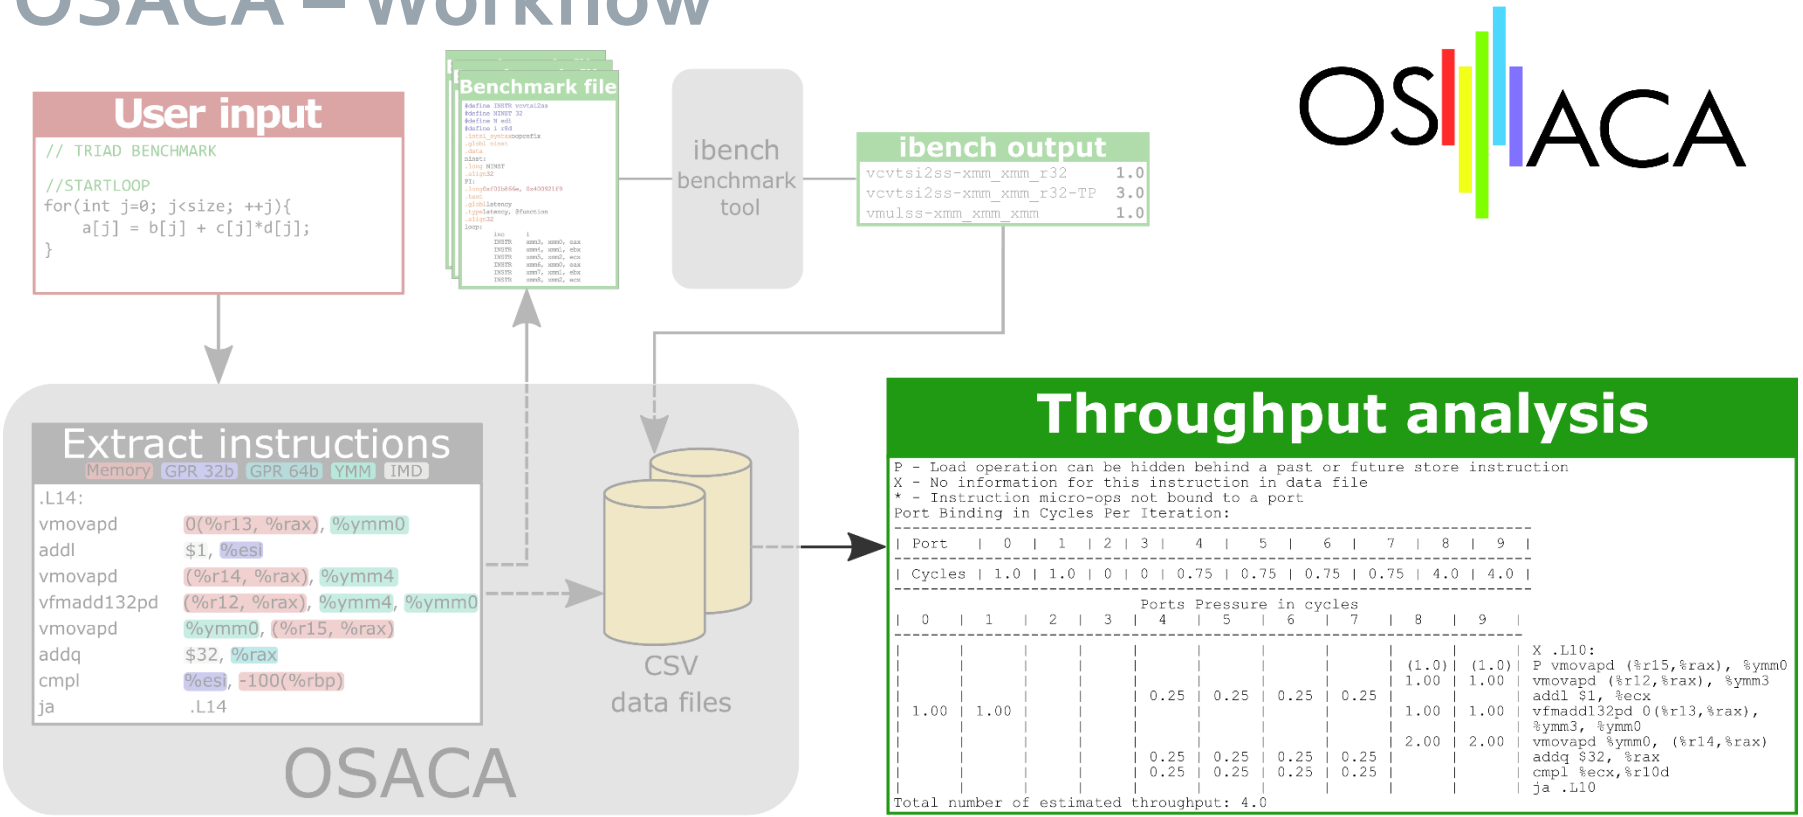

Schönauer Triad Benchmark Example

- **Load-bound**
- **Create code with -O1, -O2 and -O3 flag**
  (+ architecture specific flags)
- **Analyze for Intel Skylake & AMD Zen**

```
for(int j=0; j<size; ++j)
          a[j] = b[j] + c[j]*d[j];
```

2x unrolling

```
.L10:
  vmovaps       (%r13,%rax), %xmm0
  vmovaps       (%r15,%rax), %xmm3
  incl          %esi
  vaddpd        (%r14,%rax), %xmm3, %xmm0
  vmovaps       %xmm0, (%r12,%rax)
  addq          $16, %rax
  cmpl          %esi, %r10d
  ja            .L10
```

Insert marker for kernel detection

(done by tool or manually)

```
  movl      $111, %ebx          #START MARKER
  .byte     100, 103, 144       #START MARKER

.L10:
  vmovaps   (%r13,%rax), %xmm0
  vmovaps   (%r15,%rax), %xmm3
  incl      %esi
  vaddpd    (%r14,%rax), %xmm3, %xmm0
  vmovaps   %xmm0, (%r12,%rax)
  addq      $16, %rax
  cmpl      %esi, %r10d
  ja        .L10

  movl      $222, %ebx          #END MARKER
  .byte     100, 103, 144       #END MARKER
```

compiled for ZEN, run on ZEN

```
$ osaca --iaca --arch ZEN triad.s.zen.O3.s
$
Throughput Analysis Report
--------------------------
P - Load operation can be hidden behind a past or future store instruction
X - No information for this instruction in data file
* - Instruction micro-ops not bound to a port

Port Binding in Cycles Per Iteration:
-------------------------------------------------------------------------------------
|  Port  |   0  |   1  |   2  |   3  -DV |   4  |   5  |   6  |   7  |  8  |  9  |
-------------------------------------------------------------------------------------
| Cycles | 1.25 | 1.25 | 0.75 | 0.75   0 | 0.75 | 0.75 | 0.75 | 0.75 | 2.0 | 2.0 |
-------------------------------------------------------------------------------------

                          Ports Pressure in cycles
|  0   |  1   |  2   |  3   -  DV  |  4   |  5   |  6   |  7   |  8   |  9   |
-----------------------------------------------------------------------------
|      |      |      |             |      |      |      |      |      |      | X .L10:
| 0.25 | 0.25 | 0.25 | 0.25        |      |      |      |      | (0.5)| (0.5)| P vmovaps    0(%r13,%rax), %xmm0
| 0.25 | 0.25 | 0.25 | 0.25        |      |      |      |      | 0.50 | 0.50 | vmovaps      (%r15,%rax), %xmm3
|      |      |      |             | 0.25 | 0.25 | 0.25 | 0.25 |      |      | incl         %esi
| 0.50 | 0.50 |      |             |      |      |      |      | 0.50 | 0.50 | vaddpd       (%r14,%rax), %xmm3, %xmm0
| 0.25 | 0.25 | 0.25 | 0.25        |      |      |      |      | 1.00 | 1.00 | vmovaps      %xmm0, (%r12,%rax)
|      |      |      |             | 0.25 | 0.25 | 0.25 | 0.25 |      |      | addq         $16, %rax
|      |      |      |             | 0.25 | 0.25 | 0.25 | 0.25 |      |      | cmpl         %esi, %r10d
|      |      |      |             |      |      |      |      |      |      | ja           .L10
Total number of estimated throughput: 2.0
```

```
$ osaca --iaca --arch ZEN triad.s.zen.O3.s
$
Throughput Analysis Report
--------------------------
P - Load operation can be hidden behind a past or future store instruction
X - No information for this instruction in data file
* - Instruction micro-ops not bound to a port

Port Binding in Cycles Per Iteration:
-------------------------------------------------------------------------------------
|  Port  |   0  |   1  |   2  |   3  -DV |   4  |   5  |   6  |   7  |  8  |  9  |
-------------------------------------------------------------------------------------
| Cycles | 1.25 | 1.25 | 0.75 | 0.75   0 | 0.75 | 0.75 | 0.75 | 0.75 | 2.0 | 2.0 |
-------------------------------------------------------------------------------------

                         Ports Pressure in cycles
|  0   |  1   |  2   |  3   -  DV  |  4   |  5   |  6   |  7   |  8   |  9   |
-----------------------------------------------------------------------------
|      |      |      |             |      |      |      |      |      |      | X .L10:
| 0.25 | 0.25 | 0.25 | 0.25        |      |      |      |      | (0.5)| (0.5)| P vmovaps    0(%r13,%rax), %xmm0
| 0.25 | 0.25 | 0.25 | 0.25        |      |      |      |      | 0.50 | 0.50 | vmovaps      (%r15,%rax), %xmm3
|      |      |      |             | 0.25 | 0.25 | 0.25 | 0.25 |      |      | incl         %esi
| 0.50 | 0.50 |      |             |      |      |      |      | 0.50 | 0.50 | vaddpd       (%r14,%rax), %xmm3, %xmm0
| 0.25 | 0.25 | 0.25 | 0.25        |      |      |      |      | 1.00 | 1.00 | vmovaps      %xmm0, (%r12,%rax)
|      |      |      |             | 0.25 | 0.25 | 0.25 | 0.25 |      |      | addq         $16, %rax
|      |      |      |             | 0.25 | 0.25 | 0.25 | 0.25 |      |      | cmpl         %esi, %r10d
|      |      |      |             |      |      |      |      |      |      | ja           .L10
Total number of estimated throughput: 2.0
```

compiled for ZEN, run on ZEN

```
$ osaca --iaca --arch ZEN triad.s.zen.O3.s
$
Throughput Analysis Report
--------------------------
P - Load operation can be hidden behind a past or future store instruction
X - No information for this instruction in data file
* - Instruction micro-ops not bound to a port

Port Binding in Cycles Per Iteration:
-------------------------------------------------------------------------------------
|  Port  |   0  |   1  |   2  |   3  -DV |   4  |   5  |   6  |   7  |  8  |  9  |
-------------------------------------------------------------------------------------
| Cycles | 1.25 | 1.25 | 0.75 | 0.75   0 | 0.75 | 0.75 | 0.75 | 0.75 | 2.0 | 2.0 |
-------------------------------------------------------------------------------------

                          Ports Pressure in cycles
|  0   |  1   |  2   |  3   -  DV  |  4   |  5   |  6   |  7   |  8   |  9   |
-----------------------------------------------------------------------------
|      |      |      |             |      |      |      |      |      |      | X .L10:
| 0.25 | 0.25 | 0.25 | 0.25        |      |      |      |      | (0.5)| (0.5)| P vmovaps    0(%r13,%rax), %xmm0
| 0.25 | 0.25 | 0.25 | 0.25        |      |      |      |      | 0.50 | 0.50 | vmovaps      (%r15,%rax), %xmm3
|      |      |      |             | 0.25 | 0.25 | 0.25 | 0.25 |      |      | incl         %esi
| 0.50 | 0.50 |      |             |      |      |      |      | 0.50 | 0.50 | vaddpd       (%r14,%rax), %xmm3, %xmm0
| 0.25 | 0.25 | 0.25 | 0.25        |      |      |      |      | 1.00 | 1.00 | vmovaps      %xmm0, (%r12,%rax)
|      |      |      |             | 0.25 | 0.25 | 0.25 | 0.25 |      |      | addq         $16, %rax
|      |      |      |             | 0.25 | 0.25 | 0.25 | 0.25 |      |      | cmpl         %esi, %r10d
|      |      |      |             |      |      |      |      |      |      | ja           .L10
Total number of estimated throughput: 2.0
```

compiled for ZEN, run on ZEN

Results

Architecture executed on	Architecture compiled for	Optimization flag	Unroll factor	Measured MFLOP/s	Measured Mit/s	Measured cy/it	Prediction [cy/it] OSACA	Prediction [cy/it] IACA
Zen	Zen	**-O1**	1x	1797	898	2.00	2.00	–
Zen	Zen	**-O2**	1x	1797	898	2.00	2.00	–
Zen	**Zen**	**-O3**	**2x**	**3531**	**1754**	**1.02**	**2.00 / 2**	**–**
Skylake	Zen	**-O1**	1x	1770	885	2.03	2.00	2.24
Skylake	Zen	**-O2**	1x	1768	884	2.04	2.00	2.00
Skylake	Zen	**-O3**	**2x**	3505	1753	1.03	2.00 / 2	2.21
Zen	Skylake	**-O1**	1x	1792	896	2.01	2.00	–
Zen	Skylake	**-O2**	1x	1797	898	2.01	2.00	–
Zen	Skylake	**-O3**	**4x**	3166	1589	1.01	4.00 / 4	–
Skylake	**Skylake**	**-O1**	**1x**	**1767**	**884**	**2.04**	**2.00**	**2.24**
Skylake	Skylake	**-O2**	1x	1776	888	2.03	2.00	2.00
Skylake	Skylake	**-O3**	**4x**	6808	2738	0.53	2.00 / 4	2.21 / 4

π Benchmark Example

- **Compute-bound**

$$\pi = \int_0^1 \frac{4}{1+x^2}\,dx$$

```
int SLICES = 1000000000;
double sum = 0., delta_x = 1./SLICES;

for(int i=0; i<SLICES; ++i) {
      double x = (i+0.5)*delta_x;
      sum = sum + 4.0 / ( 1.0 + x * x);
}
double Pi = sum * delta_x;
```

8x unrolling

```
.L2:
  vextracti128 $0x1, %ymm2, %xmm0
  vcvtdq2pd    %xmm2, %ymm1
  vaddpd       %ymm7, %ymm1, %ymm1
  addl         $1, %eax
  vcvtdq2pd    %xmm0, %ymm0
  vaddpd       %ymm7, %ymm0, %ymm0
  vpaddd       %ymm8, %ymm2, %ymm2
  vmulpd       %ymm6, %ymm1, %ymm1
  vmulpd       %ymm6, %ymm0, %ymm0
  vaddpd       %ymm1, %ymm5, %ymm1
  vaddpd       %ymm0, %ymm5, %ymm0
  vdivpd       %ymm1, %ymm4, %ymm1
  vdivpd       %ymm0, %ymm4, %ymm0
  vaddpd       %ymm1, %ymm0, %ymm0
  vaddpd       %ymm0, %ymm3, %ymm3
  cmpl         $125000000, %eax
  jne          .L2
```

Compiled for SKYLAKE, run on SKYLAKE

```
$ osaca --iaca --arch SKL pi.s.skl.O3.s
-------------------------------------------------------------
|  Port  |   0  -  DV  |   1  | 2 | 3 | 4 |   5  |  6  | 7 |
-------------------------------------------------------------
| Cycles | 8.83   16.0 | 4.83 | 0 | 0 | 0 | 3.83 | 0.5 | 0 |
-------------------------------------------------------------
                  Ports Pressure in cycles
|  0   -  DV  |  1   |  2   |  3   |  4   |  5   |  6   |  7   |
---------------------------------------------------------------
|             |      |      |      |      |      |      |      | X .L2:
|             |      |      |      |      | 1.00 |      |      | vextracti128  $0x1, %ymm2, %xmm1
| 1.00        |      |      |      |      | 1.00 |      |      | vcvtdq2pd     %xmm2, %ymm0
| 0.50        | 0.50 |      |      |      |      |      |      | vaddpd        %ymm7, %ymm0, %ymm0
| 0.25        | 0.25 |      |      |      | 0.25 | 0.25 |      | addl          $1, %eax
| 1.00        |      |      |      |      | 1.00 |      |      | vcvtdq2pd     %xmm1, %ymm1
| 0.50        | 0.50 |      |      |      |      |      |      | vaddpd        %ymm7, %ymm1, %ymm1
| 0.33        | 0.33 |      |      |      | 0.33 |      |      | vpaddd        %ymm8, %ymm2, %ymm2
| 0.50        | 0.50 |      |      |      |      |      |      | vmulpd        %ymm6, %ymm0, %ymm0
| 0.50        | 0.50 |      |      |      |      |      |      | vmulpd        %ymm6, %ymm1, %ymm1
| 0.50        | 0.50 |      |      |      |      |      |      | vaddpd        %ymm0, %ymm5, %ymm0
| 0.50        | 0.50 |      |      |      |      |      |      | vaddpd        %ymm1, %ymm5, %ymm1
| 1.00   8.00 |      |      |      |      |      |      |      | vdivpd        %ymm0, %ymm4, %ymm0
| 1.00   8.00 |      |      |      |      |      |      |      | vdivpd        %ymm1, %ymm4, %ymm1
| 0.50        | 0.50 |      |      |      |      |      |      | vaddpd        %ymm1, %ymm0, %ymm0
| 0.50        | 0.50 |      |      |      |      |      |      | vaddpd        %ymm0, %ymm3, %ymm3
| 0.25        | 0.25 |      |      |      | 0.25 | 0.25 |      | cmpl          $125000000, %eax
|             |      |      |      |      |      |      |      | jne           .L2
Total number of estimated throughput: 16.0
```

Compiled for SKYLAKE, run on SKYLAKE

```
$ osaca --iaca --arch SKL pi.s.skl.O3.s
-------------------------------------------------------------
|  Port  |   0  -  DV  |   1  | 2 | 3 | 4 |   5  |  6  | 7 |
-------------------------------------------------------------
| Cycles | 8.83   16.0 | 4.83 | 0 | 0 | 0 | 3.83 | 0.5 | 0 |
-------------------------------------------------------------
                  Ports Pressure in cycles
|  0   -  DV  |  1   |  2   |  3   |  4   |  5   |  6   |  7   |
---------------------------------------------------------------
|             |      |      |      |      |      |      |      | X .L2:
|             |      |      |      |      | 1.00 |      |      | vextracti128  $0x1, %ymm2, %xmm1
| 1.00        |      |      |      |      | 1.00 |      |      | vcvtdq2pd     %xmm2, %ymm0
| 0.50        | 0.50 |      |      |      |      |      |      | vaddpd        %ymm7, %ymm0, %ymm0
| 0.25        | 0.25 |      |      |      | 0.25 | 0.25 |      | addl          $1, %eax
| 1.00        |      |      |      |      | 1.00 |      |      | vcvtdq2pd     %xmm1, %ymm1
| 0.50        | 0.50 |      |      |      |      |      |      | vaddpd        %ymm7, %ymm1, %ymm1
| 0.33        | 0.33 |      |      |      | 0.33 |      |      | vpaddd        %ymm8, %ymm2, %ymm2
| 0.50        | 0.50 |      |      |      |      |      |      | vmulpd        %ymm6, %ymm0, %ymm0
| 0.50        | 0.50 |      |      |      |      |      |      | vmulpd        %ymm6, %ymm1, %ymm1
| 0.50        | 0.50 |      |      |      |      |      |      | vaddpd        %ymm0, %ymm5, %ymm0
| 0.50        | 0.50 |      |      |      |      |      |      | vaddpd        %ymm1, %ymm5, %ymm1
| 1.00  8.00  |      |      |      |      |      |      |      | vdivpd        %ymm0, %ymm4, %ymm0
| 1.00  8.00  |      |      |      |      |      |      |      | vdivpd        %ymm1, %ymm4, %ymm1
| 0.50        | 0.50 |      |      |      |      |      |      | vaddpd        %ymm1, %ymm0, %ymm0
| 0.50        | 0.50 |      |      |      |      |      |      | vaddpd        %ymm0, %ymm3, %ymm3
| 0.25        | 0.25 |      |      |      | 0.25 | 0.25 |      | cmpl          $125000000, %eax
|             |      |      |      |      |      |      |      | jne           .L2
Total number of estimated throughput: 16.0
```

Compiled for SKYLAKE, run on SKYLAKE

```
$ osaca --iaca --arch SKL pi.s.skl.O3.s
-------------------------------------------------------------
|  Port  |   0  -  DV  |   1  | 2 | 3 | 4 |   5  |  6  | 7 |
-------------------------------------------------------------
| Cycles | 8.83   16.0 | 4.83 | 0 | 0 | 0 | 3.83 | 0.5 | 0 |
-------------------------------------------------------------
                  Ports Pressure in cycles
|  0   -  DV  |  1   |  2   |  3   |  4   |  5   |  6   |  7   |
---------------------------------------------------------------
|             |      |      |      |      |      |      |      | X .L2:
|             |      |      |      |      | 1.00 |      |      | vextracti128  $0x1, %ymm2, %xmm1
| 1.00        |      |      |      |      | 1.00 |      |      | vcvtdq2pd     %xmm2, %ymm0
| 0.50        | 0.50 |      |      |      |      |      |      | vaddpd        %ymm7, %ymm0, %ymm0
| 0.25        | 0.25 |      |      |      | 0.25 | 0.25 |      | addl          $1, %eax
| 1.00        |      |      |      |      | 1.00 |      |      | vcvtdq2pd     %xmm1, %ymm1
| 0.50        | 0.50 |      |      |      |      |      |      | vaddpd        %ymm7, %ymm1, %ymm1
| 0.33        | 0.33 |      |      |      | 0.33 |      |      | vpaddd        %ymm8, %ymm2, %ymm2
| 0.50        | 0.50 |      |      |      |      |      |      | vmulpd        %ymm6, %ymm0, %ymm0
| 0.50        | 0.50 |      |      |      |      |      |      | vmulpd        %ymm6, %ymm1, %ymm1
| 0.50        | 0.50 |      |      |      |      |      |      | vaddpd        %ymm0, %ymm5, %ymm0
| 0.50        | 0.50 |      |      |      |      |      |      | vaddpd        %ymm1, %ymm5, %ymm1
| 1.00   8.00 |      |      |      |      |      |      |      | vdivpd        %ymm0, %ymm4, %ymm0
| 1.00   8.00 |      |      |      |      |      |      |      | vdivpd        %ymm1, %ymm4, %ymm1
| 0.50        | 0.50 |      |      |      |      |      |      | vaddpd        %ymm1, %ymm0, %ymm0
| 0.50        | 0.50 |      |      |      |      |      |      | vaddpd        %ymm0, %ymm3, %ymm3
| 0.25        | 0.25 |      |      |      | 0.25 | 0.25 |      | cmpl          $125000000, %eax
|             |      |      |      |      |      |      |      | jne           .L2
Total number of estimated throughput: 16.0
```

Results

Architecture	Optimization flag	Measured MFLOP/s	Measured Mit/s	Measured cy/it	Prediction [cy/it] OSACA	Prediction [cy/it] IACA
Skylake	**-O1**	1198	200	9.02	4.75	3.91
Skylake	**-O2**	2697	450	4.00	4.25	4.00
Skylake	**-O3**	**5227**	**871**	**2.06**	**2.00**	**2.00**
Zen	**-O1**	1197	200	11.48	4.00	–
Zen	**-O2**	2696	449	4.96	4.00	–
Zen	**-O3**	5377	896	2.44	2.00	–

Compiled for SKYLAKE, run on SKYLAKE

```
$ osaca --iaca --arch SKL pi.s.skl.O1.s
$
Throughput Analysis Report
Port Binding in Cycles Per Iteration:
-------------------------------------------------------------------
|  Port  |   0  - DV  |   1  |  2  |  3  |  4  |   5  |   6  | 7 |
-------------------------------------------------------------------
| Cycles | 4.75   4.0 | 3.75 | 1.0 | 1.0 | 1.0 | 1.75 | 0.75 | 0 |
-------------------------------------------------------------------

                    Ports Pressure in cycles
|  0   -  DV  |  1   |  2   |  3   |  4   |  5   |  6   |  7   |
---------------------------------------------------------------
|             |      |      |      |      |      |      |      | X .L2:
| 0.25        | 0.25 |      |      |      | 0.25 | 0.25 |      | vxorpd       %xmm0, %xmm0, %xmm0
| 0.50        | 0.50 |      |      |      | 1.00 |      |      | vcvtsi2sd    %eax, %xmm0, %xmm0
| 0.50        | 0.50 |      |      |      |      |      |      | vaddsd       %xmm4, %xmm0, %xmm0
| 0.50        | 0.50 |      |      |      |      |      |      | vmulsd       %xmm3, %xmm0, %xmm0
| 0.50        | 0.50 |      |      |      |      |      |      | vmulsd       %xmm0, %xmm0, %xmm0
| 0.50        | 0.50 |      |      |      |      |      |      | vaddsd       %xmm2, %xmm0, %xmm0
| 1.00   4.00 |      |      |      |      |      |      |      | vdivsd       %xmm0, %xmm1, %xmm0
| 0.50        | 0.50 | 0.50 | 0.50 |      |      |      |      | vaddsd       (%rsp), %xmm0, %xmm5
|             |      | 0.50 | 0.50 | 1.00 |      |      |      | vmovsd       %xmm5, (%rsp)
| 0.25        | 0.25 |      |      |      | 0.25 | 0.25 |      | addl         $1, %eax
| 0.25        | 0.25 |      |      |      | 0.25 | 0.25 |      | cmpl         $1000000000, %eax
|             |      |      |      |      |      |      |      | jne          .L2
Total number of estimated throughput: 4.75
```

Compiled for SKYLAKE, run on SKYLAKE

```
$ osaca --iaca --arch SKL pi.s.skl.O1.s
$
Throughput Analysis Report
Port Binding in Cycles Per Iteration:
-------------------------------------------------------------------
|  Port  |   0  - DV  |   1  |  2  |  3  |  4  |   5  |   6  | 7 |
-------------------------------------------------------------------
| Cycles | 4.75   4.0 | 3.75 | 1.0 | 1.0 | 1.0 | 1.75 | 0.75 | 0 |
-------------------------------------------------------------------

                    Ports Pressure in cycles
|  0   -  DV  |  1   |  2   |  3   |  4   |  5   |  6   |  7   |
---------------------------------------------------------------------
|             |      |      |      |      |      |      |      | X .L2:
| 0.25        | 0.25 |      |      |      | 0.25 | 0.25 |      | vxorpd        %xmm0, %xmm0, %xmm0
| 0.50        | 0.50 |      |      |      | 1.00 |      |      | vcvtsi2sd     %eax, %xmm0, %xmm0
| 0.50        | 0.50 |      |      |      |      |      |      | vaddsd        %xmm4, %xmm0, %xmm0
| 0.50        | 0.50 |      |      |      |      |      |      | vmulsd        %xmm3, %xmm0, %xmm0
| 0.50        | 0.50 |      |      |      |      |      |      | vmulsd        %xmm0, %xmm0, %xmm0
| 0.50        | 0.50 |      |      |      |      |      |      | vaddsd        %xmm2, %xmm0, %xmm0
| 1.00   4.00 |      |      |      |      |      |      |      | vdivsd        %xmm0, %xmm1, %xmm0
| 0.50        | 0.50 | 0.50 | 0.50 |      |      |      |      | vaddsd        (%rsp), %xmm0, %xmm5
|             |      | 0.50 | 0.50 | 1.00 |      |      |      | vmovsd        %xmm5, (%rsp)
| 0.25        | 0.25 |      |      |      | 0.25 | 0.25 |      | addl          $1, %eax
| 0.25        | 0.25 |      |      |      | 0.25 | 0.25 |      | cmpl          $1000000000, %eax
|             |      |      |      |      |      |      |      | jne           .L2
Total number of estimated throughput: 4.75
```

Future Work

- Latency modelling
- Critical Path
- Loop-carried dependencies
- Differentiate addressing modes
- Architecture specific heuristics
- Integration in KERNCRAFT
- Replacement / Additional instrumentalization of benchmark tools

IACA – Intel Architecture Code Analyzer

Why something new?

- **OSACA is Open Source**
- **OSACA supports non-Intel architectures**
- **OSACA is based on benchmarks of individual instructions**
- **OSACA allows manual extension of the supported instruction set**
- **OSACA allows architectural exploration**

https://github.com/RRZE-HPC/OSACA

OSACA

Open Source Architecture Code Analyzer

Thank You for Your Attention!

J. Laukemann, jan.laukemann@fau.de

Department of Computer Science, FAU

Erlangen Regional Computing Center (RRZE)