# Accurate Throughput Prediction of Basic Blocks on Recent Intel Microarchitectures

Andreas Abel and Jan Reineke

UNIVERSITÄT DES SAARLANDES

SIC Saarland Informatics Campus

## Introduction

- General Problem: Predicting the throughput of compute-bound basic blocks on a specific microarchitecture.
- Existing tools (like IACA, llvm-mca, OSACA, CQA, Ithemal) are often not very accurate.
- We show that in many cases, the following simple baseline yields similar or even better predictions than existing tools:

*number of instructions* *memory reads* *memory writes*

$$max\left(1, \frac{n-1}{4}, \frac{m_r}{2}, m_w\right)$$

- We investigate how the accuracy can be improved.

## Methodology

Instruction data from uops.info

Detailed microarchitecture models obtained by reverse engineering using nanoBench

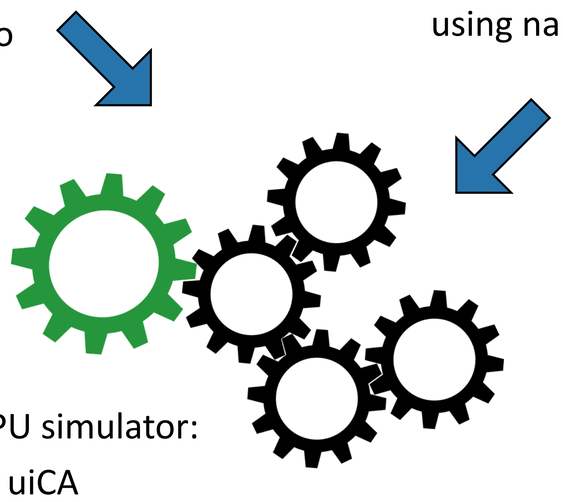

New CPU simulator: uiCA

| It. | Instruction | µops: Possible Ports | µops: Actual Port | Cycle |
|---|---|---|---|---|
| 3 | add rax, qword ptr [rcx] | {2,3} | 2 | Q I D E R |
| | | {0,1,5,6} | 5 | Q I D E R |
| | adc rcx, qword ptr [rax+rbx*1] | {2,3} | 3 | Q I D E R |
| | | {0,6} | 0 | Q I D E R |
| | shld rdx, rbx, 0x1 | {1} | 1 | Q I D E R |
| | dec r15 | {6} | 6 | Q I D E R |
| | jnz 0xffffffffffffffe1 | - | - | |
| 4 | add rax, qword ptr [rcx] | {2,3} | 2 | Q I D E R |
| | | {0,1,5,6} | 0 | Q I D E R |
| | adc rcx, qword ptr [rax+rbx*1] | {2,3} | 3 | Q I D E R |
| | | {0,6} | 0 | Q I D E R |
| | shld rdx, rbx, 0x1 | {1} | 1 | Q I D E R |
| | dec r15 | {6} | 6 | Q I D E R |
| | jnz 0xffffffffffffffe1 | - | - | |

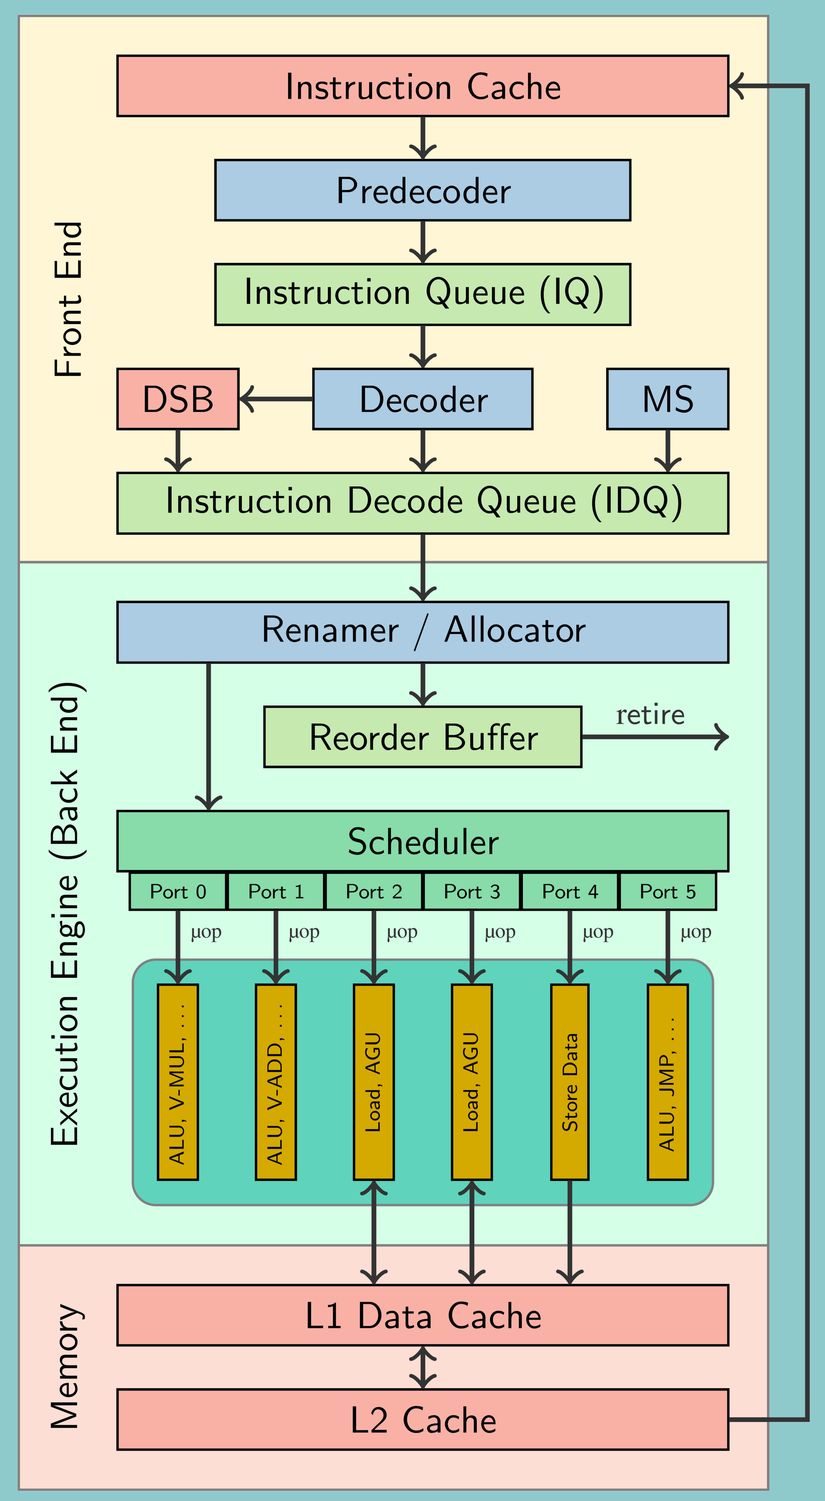


uiCA.uops.info

## Results

Comparison of the predictions of different tools for Skylake on an improved version of the BHive benchmark suite