

# Session 3: Roofline models

COMP52315: performance engineering

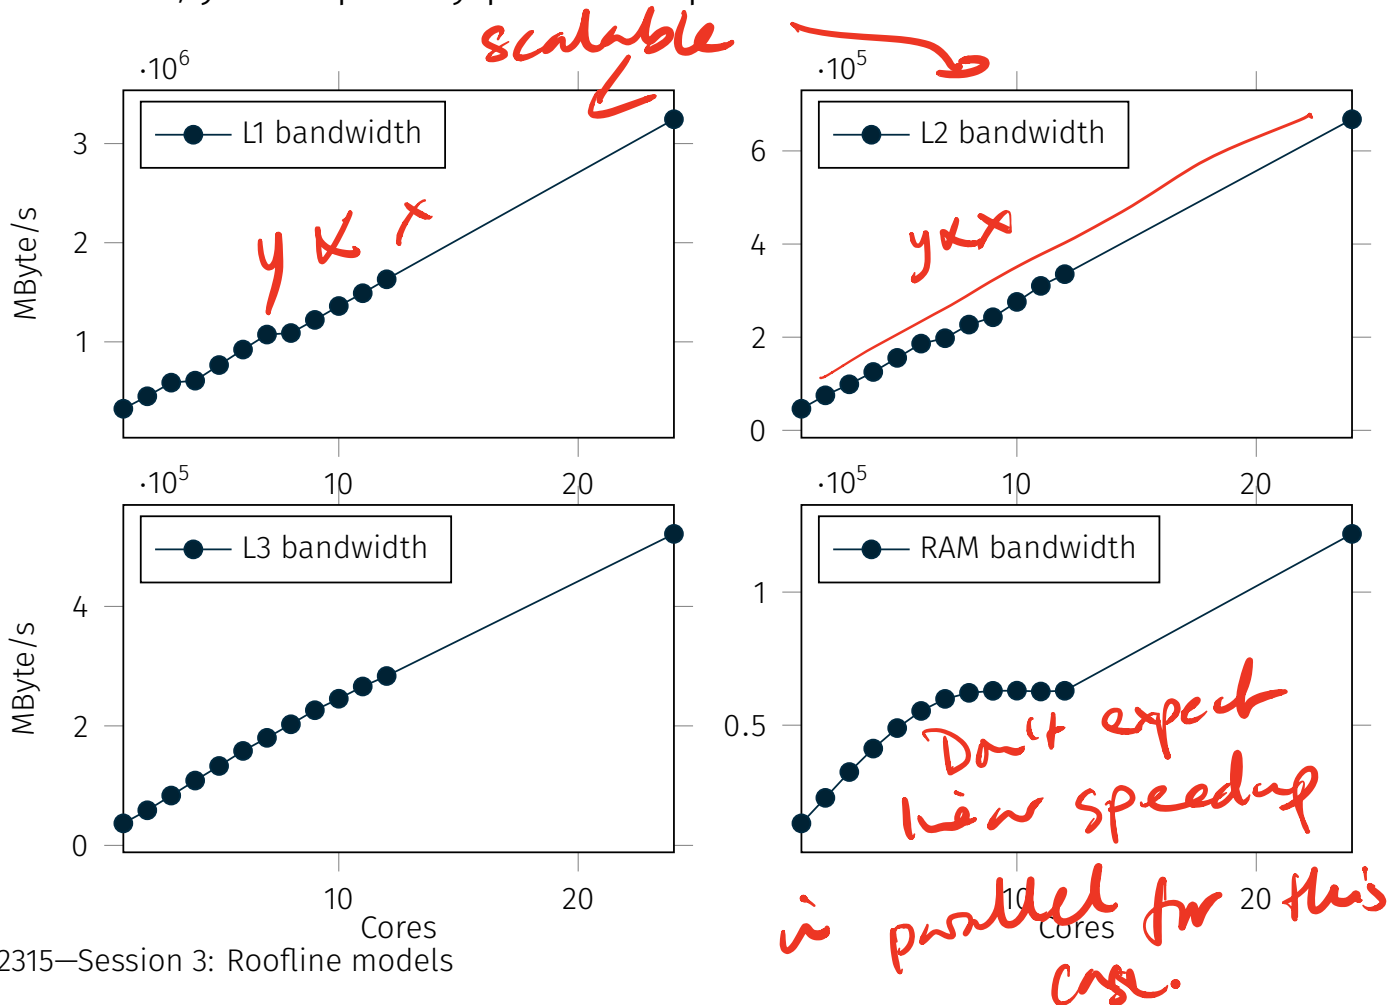
---

Lawrence Mitchell\*

\*`lawrence.mitchell@durham.ac.uk`

# Parallel memory bandwidth

In exercise 3, you hopefully produced plots similar to these.



# A more realistic measure of memory throughput

*upper bound on memory throughput.*

| The cache line copy benchmark we've seen provides upper bounds, but doesn't simulate *realistic* workloads. *likwid-bench -t-copy*

- It only touches one byte in each cache line, but remember, optimised code works on *all* the bytes in a cache line.

⇒ STREAM benchmark <https://www.cs.virginia.edu/stream/>

- Most commonly used is TRIAD.

*FORTRAN*

- Implemented in `likwid-bench` as `stream_triad_XXX` with a few different options.

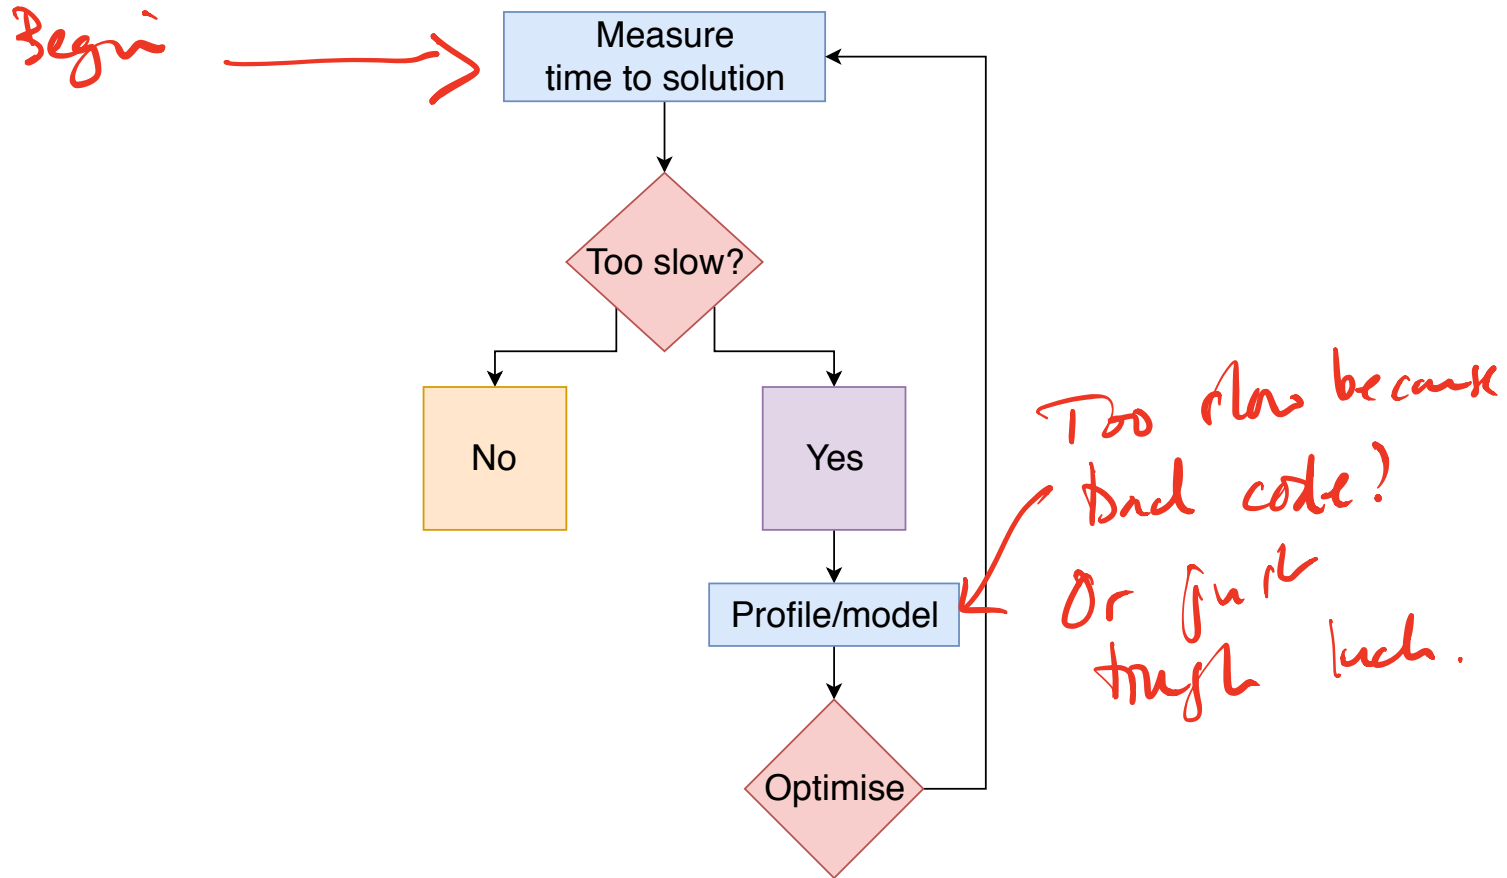
*sp, dp, sp-aux, dp-aux...*

## TRIAD loop

```
double *a, *b, *c;
double alpha = 1;
...
for (int i = 0; i < N; i++)
    a[i] = b[i]*alpha + c[i];
```


*Limited by memory throughput on all architectures & all cache levels.*

# Code optimisation



# Simple model for loop heavy code

## Simple view of hardware



## Simple view of software

```
/* Possibly nested loops */  
for (i = 0; i < ...; i++)  
  /* Complicated code doing */  
  /* N FLOPs causing  
  /* B bytes of data transfer */
```

Computational intensity [FLOPs/byte]

$$I_c = \frac{N}{B}$$

# Roofline

What is the performance  $P$  of a code?

How fast can work be done?  $P$  measured in FLOPs/s

## Bottleneck

Either

- execution of work  $P_{\text{peak}}$  [FLOPs/s];
- or the data path  $I_c b_s$  [FLOPs/byte  $\times$  byte/s].

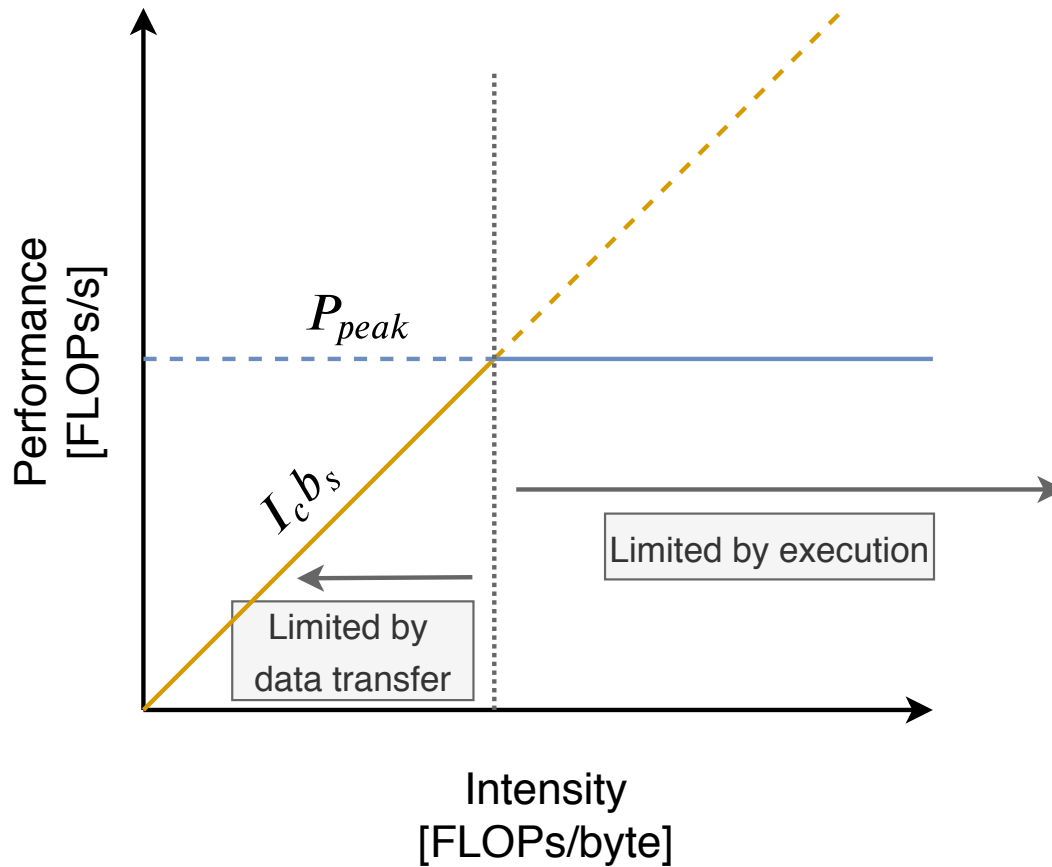
$$P = \min(P_{\text{peak}}, I_c b_s)$$

This is the simplest form of the roofline model. It is *optimistic*, everything happens at “light speed”.

Introduced in Williams et al. *Roofline: An Insightful Visual Performance Model for Multicore Architectures*, CACM (2009).

<https://doi.org/10.1145/1498765.1498785>

# Roofline



# Applying roofline

## Performance model

Roofline characterises performance using three numbers:

1.  $P_{\text{peak}}$  the peak floating point performance;
2.  $b_s$  the streaming memory bandwidth;
3.  $I_c$  the computational (or arithmetic) intensity of the code.

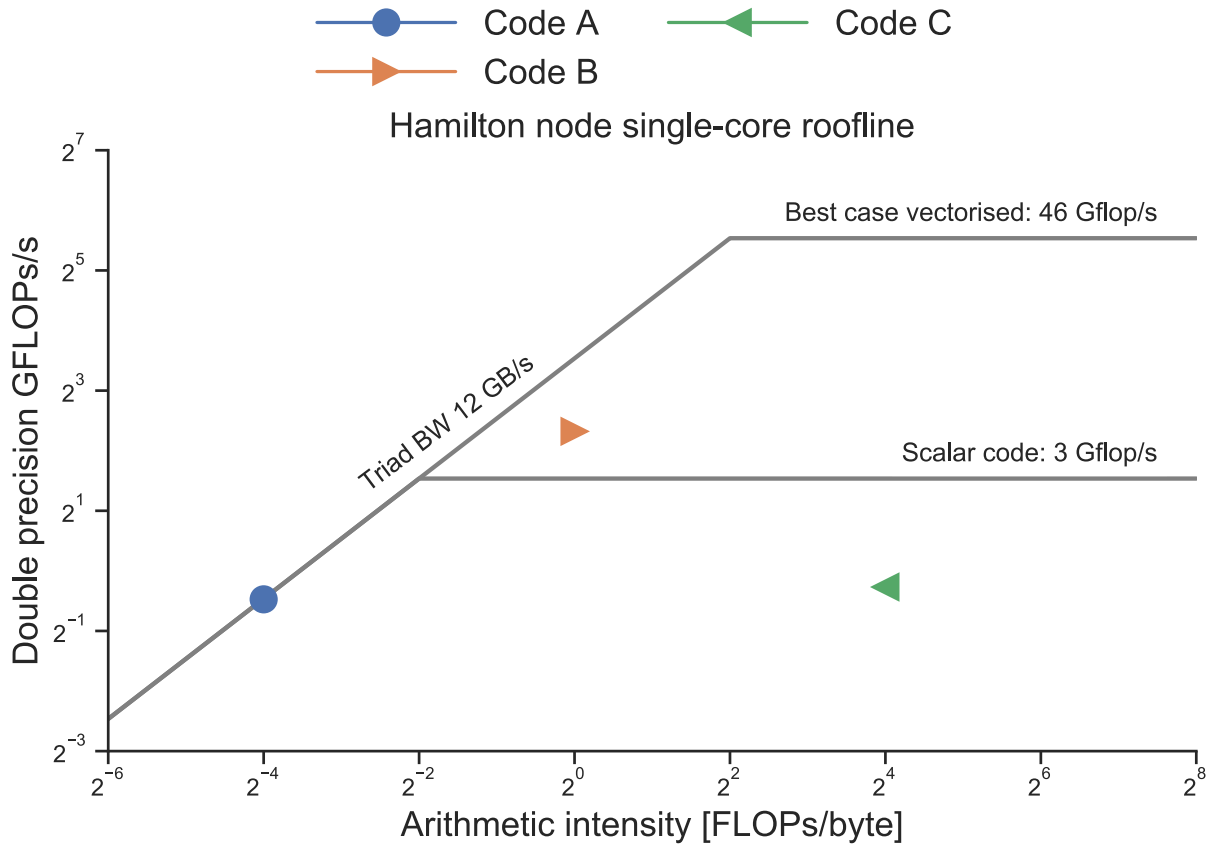
The first two are characteristics of *the hardware*. The last is a characteristic of the *code*.

## Idea


Measure these numbers and plot, gives idea of what performance optimisations are likely to pay off.



# Example



# Guides optimisation choices



Which codes might benefit from vectorisation?  
How much improvement could we expect?  
Which codes might benefit from refactoring to increase arithmetic intensity?

## Memory bandwidth

Roofline models data movement with streaming memory bandwidth.

Two ways of computing it.

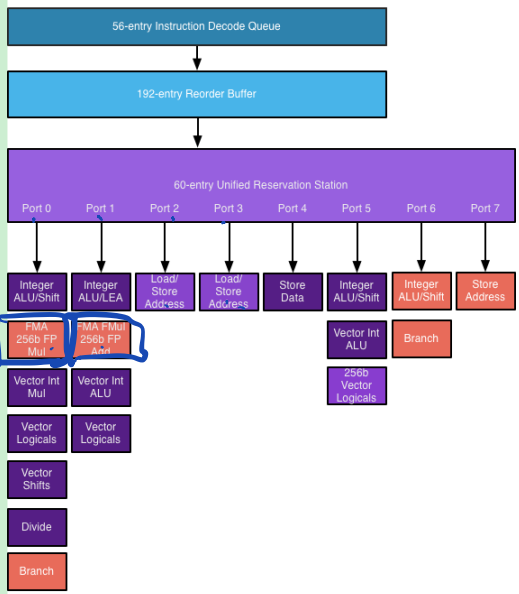
1. Know what speed of memory you have, and look up number of memory channels on spec sheet. For example, 4-channel 2.4GHz RAM delivers at best  $4 \times 2.4\text{GHz} \times 8\text{Byte} = 76.8\text{GByte/s}$ .  
⇒ Needs knowledge of installed memory, typically not achieved in practice.
2. *Measure* using STREAM.  
⇒ we will typically do this (see exercise 4).

# Determining machine characteristics

## Floating point throughput

Absolute peak can be determined from spec sheet frequency and some knowledge of hardware.

### Intel Haswell Execution Engine




- Floating point instructions execute on port 0 and port 1
- Up to 4 “micro-ops” issued per cycle ⇒ up to 2 floating point instructions per cycle
- FMA ( $y \leftarrow a + b \times c$ ); MUL execute on both ports.
- ADD only executes on port 1. Divide only executes on port 0.

# Determining machine characteristics

## Example: best case

Code only contains double precision SIMD FMAs, clock speed is 2.9GHz.

Peak floating point throughput is



$$\underbrace{2.9}_{\text{clock speed}} \times \underbrace{2}_{\text{dual issue}} \times \underbrace{4}_{\text{vector width}} \times \underbrace{2}_{\text{FMA}} = \underline{46.4 \text{ GFLOPs/s}}$$

## Example: only ADDs

Code only does double precision SIMD ADDs, clock speed is 2.9GHz.



$$\underbrace{2.9}_{\text{clock speed}} \times \underbrace{1}_{\text{single issue}} \times \underbrace{4}_{\text{vector width}} = 11.6 \text{ GFLOPs/s}$$

# Determining machine characteristics

Dense matrix solve  $X = \underline{A^{-1}B}$

- Often useful to put multiple “roofs” on the roofline, corresponding to different instruction mixes.
  - Calculations are complicated by frequency scaling as well.
- ⇒ can add measured limit by running LINPACK (see exercises)

## More details

<https://uops.info> has all the information you could ever want on micro-op execution throughput.

<https://travisdowns.github.io/blog/2019/06/11/speed-limits.html> discusses in much more detail how to find limiting factors in (simple) code.

# Computing arithmetic intensity

Two options:

1. Measure using performance counters (see later);
2. Read code, count floating point operations and data accesses.

Both options have their pros and cons.

# Counting operations

```
double *a, *b, *c, *d;  
...  
for (i = 0; i < N; i++) {  
    a[i] = b[i]*c[i] + d[i]*a[i];  
}
```

1      3      2

2	MUL
1	ADD

3 DP FLOPs/iteration.  $3N$  total DP FLOPs. (Notice how we don't care about what type of FLOPs these are).



# Counting data accesses

Each read counts as one access. Each write counts as two (one load, one store). Only care about array data (ignore loop variables)

```
double *a, *b, *c, *d;
...
for (i = 0; i < N; i++) {
    a[i] = b[i]*c[i] + d[i]*a[i];
}
```

*Handwritten annotations:* Under the code, there are blue numbers: '4, 5' under the closing brace, '1' under 'b[i]', '2' under 'c[i]', '3' under 'd[i]', and '4' under 'a[i]'. There is also a blue 'α' above the '+' sign.

*Handwritten notes:* Write first loads data or then stores it  
1 load  
+ 1 store.

3 DP reads, 1 DP write per iteration. 8 × 5N total bytes.

~~OpStu~~  $OI = \frac{3}{40}$

# Complication

```
double *a, *b, *c, *d;
...
for (i = 0; i < N; i++)
    for (j = 0; j < M; j++)
        a[j] = b[i]*c[i] + d[i]*a[j];
```

For actual data moved, need a *model* of cache.

## Bounds on movement

### Perfect cache

Provides lower bound.

Each array entry moved from main memory once.  $8(3N + 2M)$

Counts *unique* memory accesses.  $3N$  reads +  $2M$  writes

COMP23152:  $8 \times 2MN + 8 \times 3MN$  total bytes.

### Pessimal cache

Provides upper bound

Each array access misses cache.

Counts *total non-unique* memory accesses.

$8 \times 2MN + 8 \times 3MN$  total bytes.

# Complication

## Bounds on movement

### Perfect cache

Provides lower bound.

Each array entry moved from main memory once.

Counts *unique* memory accesses.

$8 \times 2M + 8 \times 3N$  total bytes.

### Pessimal cache

Provides upper bound

Each array access misses cache.

Counts *total non-unique* memory accesses.

$8 \times 2MN + 8 \times 3MN$  total bytes.

These bounds are typically not tight. If you want better bounds normally have to work harder in the analysis.

Best employed in combination with measurement of arithmetic intensity.

# Exercise: roofline plot for dense matrix-vector multiplication

- Goal is to produce a roofline plot for dense matrix-vector multiplication, which computes

$$\vec{y} = A\vec{x} = \sum_j A_{ij}\vec{x}_j$$

https:

`//teaching.wence.uk/comp52315/exercises/exercise04/`