P. Hadjidoukas

Set 6 - Amdahl's Law,Roofline Model

Issued: April 26, 2023

Question 1: Amdahl's Law

a) Assume you work on Euler and you have one node with 24 cores that you can use to solve a problem in parallel for which 91% of your code is parallelizable. Can you get a speedup of 8? If so, how many cores are needed at least?

This problem can be solved with Amdahl's law formula: If we solve for X, we find that we need 26 cores. The code offers not enough parallelism to achieve a speedup of 8 on one node of Euler.

$$S(p) = \frac{T(1)}{T(p)} = \frac{T(1)}{0.09T(1) + 0.91T(1)/p} = \frac{1}{0.09 + 0.91/p} \overset{!}{=} 8$$

$$\Rightarrow p = \frac{8 \times 0.91}{1 - 8 \times 0.09} = 26$$

b) Profiling a serial code for Molecular Dynamics you find that 90% of the time is spent in a large loop with independent iterations (perfectly parallelizable with N threads), another 5% is spent in a region that can be parallelized with at most 2 threads and the remaining part is purely serial.
Given Amdahl's law, what is the strong scaling for $N \to \infty$?
For what value of N is the speedup equivalent to 90% of the asymptotic maximum?

$$\frac{T_1}{T_N} = \frac{1}{0.05 + \frac{0.05}{2} + \frac{0.90}{N}} = \frac{200N}{15N + 180} \quad (1)$$

$$\lim_{N\to\infty} \frac{T_1}{T_N} = \lim_{N\to\infty} \frac{200N}{15N + 180} = \frac{200}{15} = \frac{40}{3} \approx 13.33 \quad (2)$$

$$\frac{200N}{15N + 180} = \frac{90}{100}\frac{40}{3} = 12 \Rightarrow N = 108 \quad (3)$$

Question 2: Roofline Model

Given the following serial code snippet:

```
float A[N], B[N], C[N];
...
const int P=2;
for (int i=0; i<N; ++i ) {
  int j=0 ;
  while (j < P) {
     A[i] = B[i]*A[i]+0.5 ;
    ++j;
  }
  C[i] = 0.9*A[i]+C[i] ;
}
```

a) What is the operational intensity of the code? Assume an infinite cache and state any further assumption you made. Show your calculations.

Assumptions: we have enough registers.

```
A[i] = B[i]*A[i] + 0.5;
A[i] = B[i]*A[i] + 0.5;
C[i] = 0.9*A[i]+ C[i];
```

memory operations: read A,B, C write A, C, (3+2)*4=20 bytes, fp operations: 3 mul + 3 add = 6 flops

$OI = 6/20$

Make sure that integer operations (++j) are not counted.

b) A compute node has a peak performance of 409.7 GFLOP/s (single precision) and a memory bandwidth of 34 GB/s. For which range of positive integer values P is the code of subquestion (a) memory bound? Show your calculations.

Ridge point = 409.7 / 34 = 12.05 Assuming perfect caching meaning we have enough registers so A[i], B[i], C[i] are read only once from the memory and A[i], C[i] are written only once to the memory.

```
A[i] = B[i]*A[i] + 0.5; // 2 FLOP
 ...
A[i] = B[i]*A[i] + 0.5;  // 2 FLOP
C[i] = 0.9*A[i] + C[i];  // 2 FLOP
```

$OI < \text{Ridge point} = 12.05$

Flops = (2 + 2P), Bytes = 5*4 = 20

$(2 + 2P)/20 < 12.05$

$2 + 2P < 20 * 12.05 = 241$

$2P < 239$

$P < 239/2 = 119.5$

$P <= 119$