

Parallel Programming: Speedups and Amdahl's law

Oregon State University
Mike Bailey
mjb@cs.oregonstate.edu

CC BY-NC-ND 4.0 International License

speedups and amdahl's law.pdf mjb - March 21, 2021

1

Definition of Speedup

If you are using n processors, your **Speedup_n** is:

$$Speedup_n = \frac{T_1}{T_n}$$

where T_1 is the execution time on **one core** and T_n is the execution time on **n cores**. Note that Speedup_n should be > 1 .

And your **Speedup Efficiency_n** is:

$$Efficiency_n = \frac{Speedup_n}{n}$$

which could be as high as 1., but probably never will be.

Oregon State University Computer Graphics mjb - March 21, 2021

2

However, Multicore is not a Free Lunch: Amdahl's Law

If you put in n processors, you should get n times Speedup (and 100% Speedup Efficiency), right? Wrong!

There are always some fraction of the total operation that is inherently sequential and cannot be parallelized no matter what you do. This includes reading data, setting up calculations, control logic, storing results, etc.

If you think of all the operations that a program needs to do as being divided between a fraction that is parallelizable and a fraction that isn't (i.e., is stuck at being sequential), then **Amdahl's Law** says:

$$Speedup_n = \frac{T_1}{T_n} = \frac{1}{\frac{F_{parallel}}{n} + F_{sequential}} = \frac{1}{\frac{F_{parallel}}{n} + (1 - F_{parallel})}$$

This fraction can be reduced by deploying multiple processors. This fraction can't.

Oregon State University Computer Graphics mjb - March 21, 2021

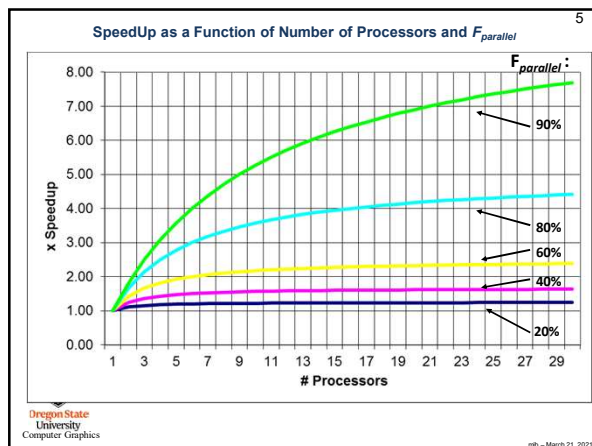
3

A Visual Explanation of Amdahl's Law


The Sequential Portion doesn't go away, and it also doesn't get any smaller. It just gets more and more dominant.

Oregon State University Computer Graphics mjb - March 21, 2021


4




5



6



7



8

You can also solve for $F_{parallel}$ using Amdahl's Law if you know your speedup and the number of processors

Amdahl's law says:

$$S = \frac{T_s}{T_p} = \frac{1}{\frac{F}{n} + (1-F)} \Rightarrow \frac{1}{S} = \frac{F}{n} + (1-F) = 1 + \frac{F-nF}{n} \Rightarrow \frac{1}{S} - 1 = F \frac{(1-n)}{n}$$

Solving for F:

$$F = \frac{\frac{1}{S} - 1}{\frac{1-n}{n}} = \frac{\frac{T_s}{T_p} - 1}{\frac{1-n}{n}} = \frac{\frac{T_s - T_p}{T_p}}{\frac{1-n}{n}} = \frac{n(T_s - T_p)}{T_p(1-n)} = \frac{n(T_1 - T_n)}{T_1(n-1)}$$

Use this if you know the timing

Use this if you know the speedup

$$F = \frac{n}{(n-1)} \frac{T_1 - T_n}{T_1} = \frac{n}{(n-1)} \left(1 - \frac{1}{Speedup} \right)$$

If you've got several (n,S) values, you can take the average (which is actually a least squares fit):

$$\bar{F}_i = \frac{n_i}{(n_i - 1)} \frac{T_1 - T_{n_i}}{T_1}, i = 2..N$$

$$\bar{F} = \frac{\sum_{i=2}^N F_i}{N-1}$$

note that when $i=1$, $T_n = T_1$

Oregon State University Computer Graphics

mp - March 21, 2021

9

Amdahl's Law can also give us the Maximum Possible SpeedUp

Note that these fractions put an upper bound on how much benefit you will get from adding more processors:

$$\max_{n \rightarrow \infty} Speedup = \lim_{n \rightarrow \infty} Speedup = \frac{1}{F_{sequential}} = \frac{1}{1 - F_{parallel}}$$

$F_{parallel}$	maxSpeedup
0.00	1.00
0.10	1.11
0.20	1.25
0.30	1.43
0.40	1.67
0.50	2.00
0.60	2.50
0.70	3.33
0.80	5.00
0.90	10.00
0.95	20.00
0.99	100.00

Oregon State University Computer Graphics

mp - March 21, 2021

10

A More Optimistic Take on Amdahl's Law: The Gustafson-Baris Observation

Gustafson observed that as you increase the number of processors, you have a tendency to attack larger and larger versions of the problem. He also observed that when you use the same parallel program on larger datasets, the parallel fraction, F_p , increases.

Let P be the amount of time spent on the parallel portion of an original task and S spent on the serial portion. Then

$$F_p = \frac{P}{P+S} \quad \text{or} \quad S = \frac{P - PF_p}{F_p}$$

Without loss of generality, we can set $P=1$ so that, really, S is now a fraction of P . We now have:

$$S = \frac{1 - F_p}{F_p}$$

Oregon State University Computer Graphics

mp - March 21, 2021

11

A More Optimistic Take on Amdahl's Law: The Gustafson-Baris Observation

We know that if we multiply the amount of data to process by N , then the amount of parallel work becomes NP . Surely the serial work must increase too, but we don't know how much. Let's say it doesn't increase at all, so that we know we are getting an upper bound answer.

In that case, the new parallel fraction is: $F'_p = \frac{P'}{P'+S} = \frac{NP}{NP+S}$

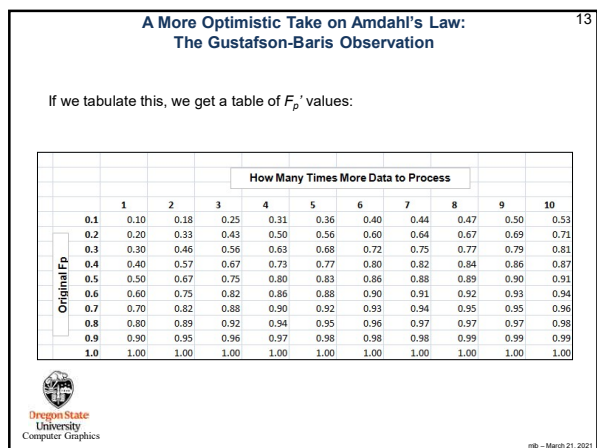
And substituting for $P (=1)$ and for S , we have:

$$F'_p = \frac{N}{N+S} = \frac{N}{N + \frac{1-F_p}{F_p}}$$


Oregon State University Computer Graphics

mp - March 21, 2021


12



13



14



15