

PSC II: More computational scaling

Last time saw

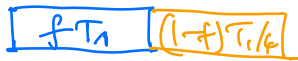
Strong scaling (Amdahl) 1960s

$$T_p = f T_1^{\text{global}} + \frac{(1-f) T_1^{\text{global}}}{p}$$

Fixed global problem size, add more resource \rightarrow



get answer faster.

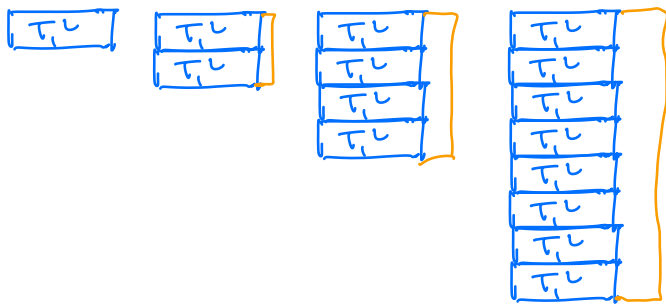


Modern supercomputers have $O(10^6)$ cores.

Weak scaling (Gustafson) 1990s

$$T_p = T_1^{\text{local}} + \underbrace{O(p) T_1^{\text{local}}}_{\text{overhead}}$$

Fixed local problem size, add more resource \rightarrow solve bigger problem in same time.



Relative time on one process is $p T_1^{\text{local}}$

Speedup (Amdahl)

$$S_p = \frac{T_1}{T_p} = \frac{T_1}{fT_1 - \frac{(1-f)T_1}{P}} = \frac{1}{f - \frac{1-f}{P}}$$
$$= \frac{P}{(P+1)f - 1}$$

$f = 0.1 \quad P = 100$
 $S_p = \frac{100}{10 \cdot 1 - 1} \approx 11$

Speedup (Gustafson)

$$S_p = \frac{P T_i^{\text{local}}}{T_i^{\text{local}} + o(p) T_i^{\text{local}}} = \frac{P}{1 + o(p)}$$



Efficiency

Amdahl

$$\eta_p = \frac{T_1}{P T_p} = \frac{1}{(P+1)f - 1}$$

cost on one process

cost on p processes

Gustafson

$$\eta_p = \frac{P T_i^{\text{local}}}{P(T_i^{\text{local}} + o(p) T_i^{\text{local}})} = \frac{1}{1 + o(p)}$$

Why is timestepping a strong scaling problem?

Parabolic problems (eg heat equation)

$$\frac{\partial u}{\partial t} - \underbrace{\nabla^2 u}_{\text{diffuses out}} = f \quad \text{external forcing}$$


Stable explicit timestep is $O(h^2)$ for grid spacing h .

$$\frac{u^{n+1} - u^n}{\Delta t} - \nabla^2 u^n = f \quad \text{explicit Euler}$$

\xrightarrow{h}

$$\Delta t = O(h^2)$$

\rightarrow This is why explicit schemes are terrible for parabolic problems.

Hyperbolic problems (eg advection equation)

$$\frac{\partial \phi}{\partial t} + \nabla \cdot (\phi \underline{u}) = f$$

\uparrow fluid velocity

Stable explicit timestep is $O(h)$.

What's the consequence for parallel computing?

Discrete heat equation

\Rightarrow Grid spacing $h \Rightarrow$ answer
in 1 hr. On one process.

But now want spacing $\frac{h}{2}$

\Rightarrow In 2D this is 4x work per
timestep.

OK: sounds good. 4x work / timestep.

But problem is 4x bigger
so \Rightarrow use 4x processes.

Problem for stability I also need
4x timesteps.

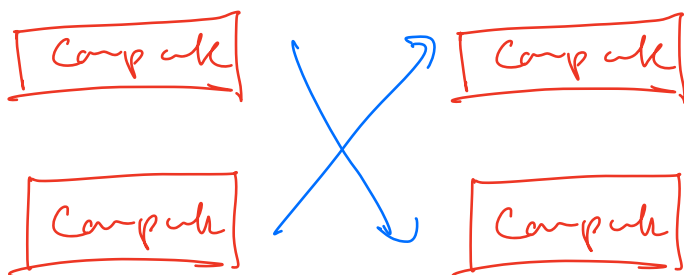
So to get the answer in 1 hr,
need to do 4x as many
timesteps.

\Rightarrow add more compute

\rightarrow local problem gets smaller
(strong scaling)

What are sources of serial fraction and parallel overhead?

- Communication latency



message exchange phase

→ we'll reuse this

Time to send b-byte message

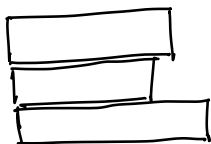
$$t(b) = \alpha + \beta b$$

latency

inverse bandwidth

Rich hardware designers.

- Load imbalance



Each process takes a different amount of time to do local work.

⇒ perhaps # of parts it sees is different.

This creates a "serial fraction".

⇒ for multigrid this can happen early

Grid-based PDEs: design considerations

- Nothing that is $O(P)$.

Typically you might have a lookup array that tells you which process a dof belongs to.

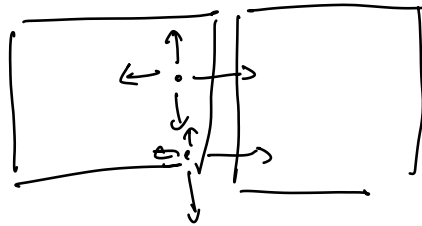
\Rightarrow we can never "gather" all the data to a single process.

- At worst $O(\log P)$ communication complexity.

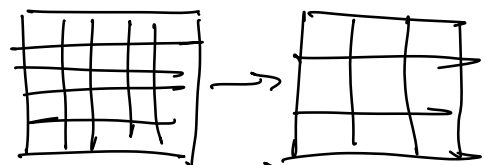
\rightarrow Nearest neighbour

\rightarrow Multilevel

Stencil
 $\sim P$



Stencils only access a neighborhood ✓



coarsening

gives us $\log P$ behaviour for long range interactions

- At worst $O(N \log N)$ compute complexity.

multigrid gives us this, coarsening factor

$$\sum_{i=1}^{\infty} \frac{1}{2^i} < 2.$$

Machine characteristics

Messaging:

$$t(d) = \alpha + \beta d$$

↑
latency

inverse bandwidth



βd

⇒ write ping pong code.