

Case study: adding n numbers

Cost = $pT_p \propto K.T_s$ for cost-efficiency

$$T_0 = pT_p - T_s$$

$p \downarrow$ Cases \downarrow	T_s	T_p	Cost pT_p	Speedup $\frac{T_s}{T_p}$	Efficiency $\frac{T_s}{p.T_p}$
(A) $p=1$	n	n	n	1	1
(B) $p=n$	n	$\log n$	$n \log n$	$\frac{n}{\log n}$	$\frac{1}{\log n}$
(C) $p < n$ blind map	n	$\frac{n}{p} \log n$	$n \log n$	$\frac{p}{\log n}$	$\frac{1}{\log n}$
(D) $p < n$ $x \rightarrow x \bmod p$	n	$\frac{n}{p} \log p + \frac{n}{p}$ $= O\left(\frac{n}{p} \log p\right)$	$n \log p$	$\frac{p}{\log p}$	$\frac{1}{\log p}$
(E) $p < n$ $x \rightarrow \lfloor \frac{x}{p} \rfloor$	n	$\log p + \frac{n}{p}$	$n + p \log p$	$\frac{np}{p \log p + n}$	$\frac{n}{p \log p + n}$

\hookrightarrow if $\theta(n) \stackrel{?}{=} \theta(n + p \log p)$
 $\theta(n) = \theta(p \log p)$

PERFORMANCE & SCALABILITY

• parallel system \equiv algorithm + architecture


• Speedup $S = \frac{T_s}{T_p}$ \rightarrow ?

• Superlinear speedup?


• Efficiency $E = S/p = \frac{T_s}{p \cdot T_p}$ • Cost = $p \cdot T_p$

• Cost-optimal if


Cost $\propto T_s$,
 $E = \Theta(1)$




(a) Initial data distribution and the first communication step




(b) Second communication step



(c) Third communication step



(d) Fourth communication step




(e) Accumulation of the sum at processor 0 after the final communication

Figure 4.1 Computing the sum of 16 numbers on a 16-processor hypercube. Σ_i^j denotes the sum of numbers with consecutive labels from i to j .
 Copyright (r) 1994 Benjamin/Cummings Publishing Co.


• $S = \Theta\left(\frac{n}{\log n}\right)$ • $E = \Theta\left(\frac{1}{\log n}\right)$ • Cost = $\Theta(n \log n)$
 • not cost optimal

• Examine effect of granularity and data mapping:
 "if system with n procs is cost-optimal, using p procs ($p < n$) to simulate n procs preserves cost-optimality"

- $T_p = \# \text{transfers} \times \# \text{steps}$
- $\Theta\left(\frac{n}{p} \log p\right) + \Theta\left(\frac{n}{p}\right) = \Theta\left(\frac{n}{p} \log p\right) \therefore \text{not cost-optimal}$
 - Overhead $T_0 = \Theta(n \log p)$
 - Problem size w to add n numbers is $\Theta(n)$




(a) Four processors simulating the first communication step of 16 processors



(b) Four processors simulating the second communication step of 16 processors

Figure 4.2 Four processors simulating 16 processors to compute the sum of 16 numbers (first two steps). Σ_i^j denotes the sum of numbers with consecutive labels from i to j .
Copyright (r) 1994 Benjamin/Cummings Publishing Co.

• Isoefficiency function does not exist because $[w = KT_0(w, p)]$ cannot be satisfied for any $K \therefore$ system not scalable



(c) Simulation of the third step in two substeps



(d) Simulation of the fourth step

(e) Final result

Figure 4.3 (cont.) Four processors simulating 16 processors to compute the sum of 16 numbers (last three steps).

Copyright (r) 1994 Benjamin/Cummings Publishing Co.

• $\Theta(n/p) + \Theta(\log p) = T_p$

• Cost = $\Theta(n + p \log p)$ = cost-optimal as long as $n = \Omega(p \log p)$

asymptotics

• $T_p = \frac{n}{p} - 1 + 2 \log p$


• $S = \frac{n}{n/p + 2 \log p} = \frac{np}{n + 2p \log p}$ constants

• $E = \frac{n}{n + 2p \log p}$


• $T_0(w, p) = p \left(\frac{n}{p} + 2 \log p \right) - n = 2p \log p$

• (isoefficiency func): ~~W~~ $W = 2K p \log p$ i.e., $\Theta(p \log p)$


• $W \approx n$ & $T_0 = \Theta(p \log p)$ & condn. for cost-optimality: $W = \Omega(p \log p)$



(a)



(b)



(c)



(d)

Figure 4.4 A cost-optimal way of computing the sum of 16 numbers on a four-processor hypercube.
Copyright (r) 1994 Benjamin/Cummings Publishing Co.

$W \propto p \log p$
 $W' \propto p' \log p'$
 \therefore increasing p by factor $\frac{p'}{p}$ requires increasing workload by factor $\frac{p' \log p'}{p \log p}$.
to increase speedup by factor p'/p

$E = \frac{T_s}{p T_p}$; $T_0 = p T_p - T_s \therefore E = \frac{1}{1 + \frac{T_0}{T_s}}$

T_0 : \uparrow fn of p . fgm has serial component, communication, idling etc. (T_{serial})

$T_0 \propto (p-1) T_{serial}$

\therefore For fixed problem size (ie, fixed T_s), as $p \uparrow$, $T_0 \uparrow$ and $E \downarrow$


$\uparrow n$, keeping p constant: In many cases, T_0 grows sublinearly w.r.t. T_s

$\Rightarrow E \uparrow$

From (A) & (B), possible to keep efficiency fixed, by $\uparrow n$ & $\uparrow p$ --- "scalable"

Q] Rate at which problem size $n \uparrow$ w.r.t. p , to keep E fixed?

- ① Speedup does not \uparrow linearly as #procs increases $\therefore E \downarrow$
 - ② larger instance of same problem yields $\uparrow S$ & E (for fixed p)
- Very common trends



- Scalable parallel system $\equiv (\uparrow p \text{ \& \; problem size})$ to keep E fixed
- Scalability \equiv measure of capacity to $\uparrow S$ in proportion to p
- $E = \Theta(1)$ for cost-optimal system
- Scalability & cost-optimality are related: scalable system can always be made cost-optimal if p and (size of computation) chosen correctly.

eg below:
 system stays cost-optimal at $E=0.8$ if n is increased as $p \log p$




Figure 4.5 Speedup versus the number of processors for adding a list of numbers on a hypercube.
 Copyright (r) 1994 Benjamin/Cummings Publishing Co.

n	$p=1$	$p=4$	$p=8$	$p=16$	$p=32$
64	1	0.8	0.57	0.33	0.17
192	1	0.92	0.8	0.6	0.38
320	1	0.95	0.87	0.71	0.5
512	1	0.97	0.91	0.8	0.62

[Efficiency as a function of n & p for Fig 4.4]

• Cost-optimal when $n = \sqrt{2} (p \log p)$. For $\begin{cases} n=64 \\ p=4 \end{cases}$, $n = 8 p \log p$

$64 = 8 \times [4 \log 4]$ $512 = 8 [16 \log 16]$

$192 = 8 \times [8 \log 8]$

- performance of scaled down algo may be different for different assignments of (virtual \rightarrow real) processors

eg $[n \times n] \times [n \times 1]$ on p -proc HC $\begin{cases} \rightarrow p \text{ square blocks} \\ \rightarrow p \text{ stripes of } n/p \text{ cols/rows} \end{cases}$

ISOEFFICIENCY METRIC OF SCALABILITY

- useful to determine rate at which problem size must \uparrow w.r.t. p to keep the efficiency fixed

Defn:

- Problem size $W = \#$ computation steps in best seq. algo. on 1 proc.

- Overhead $T_0(W, p) = pT_p - W$

$$T_p = (T_0 + W) / p$$

$\parallel W = T_0$
 $\parallel T_0$ is \uparrow ing function of p

$$S = W / T_p = \frac{Wp}{T_0 + W}$$

$$E = \frac{S}{p} = \frac{W}{W + T_0} = \frac{1}{1 + T_0(W, p) / W} \quad \cdot \quad \text{Examine effect on } E!$$

Typically, T_0 grows slower than $\Theta(W)$ for a fixed p \parallel observation

$$\underline{W = \left(\frac{E}{1-E} \right) T_0(W, p) = \underline{KT_0(W, p)} \quad \text{----- isoefficiency func.}}$$

- func determines ease with which parallel system can maintain const. efficiency & hence achieve speedups increasing in proportion to p
- small \Rightarrow small \uparrow in W sufficient to efficiently utilize \uparrow # proc
- large \Rightarrow poorly scalable
- T_0 may have terms of different orders of magnitude component that requires W to grow at the highest rate w.r.t. p determines overall asymptotic isoefficiency function

eg $T_0 = p^{3/2} + p^{3/4} W^{3/4}$ $\begin{cases} \text{--- } W = K p^{3/2} \\ \text{--- } W = K p^{3/4} W^{3/4} \Rightarrow W = K^4 p^3 \end{cases}$

$\therefore \Theta(p^{3/2})$ 1st term & $\Theta(p^3)$ 2nd term

- Parallel system is cost-optimal iff $p \cdot T_p = \Theta(W)$
i.e., $W + T_0(W, p) = \Theta(W)$ // $T_0 = pT_p - W$
 $T_0(W, p) = O(W)$
 $W = \Omega(T_0(W, p))$ } i.e. iff T_0 does not asymptotically exceed $\frac{W}{p}$
- if $W = K T_0(W, p)$ gives isoefficiency function $f(p)$, then $W = \Omega(f(p))$ must hold to ensure cost-optimality with scaling

LOWER BOUND ON ISOEFFICIENCY FUNCTION

- Asymptotically, W must increase at least as fast as $\Theta(p)$ to maintain fixed efficiency
 $\therefore \Omega(p)$ is asymptotic lower bound on isoefficiency fn.
 Ideally, $\Theta(p)$.

DEGREE OF CONCURRENCY & ISOEFFICIENCY FUNCTION

- measure of # eps in parallel as a function of W
- eg solve n eqns in n variables using Gaussian elimination
 Computation = $\Theta(n^3)$ totally, but n vars are sequentially eliminated, each needs $\Theta(n^2)$ computations
 $W = \Theta(n^3)$, $C(W) = \Theta(W^{2/3})$
 Given p processors, problem size $W \geq \Omega(p^{3/2})$ to use all procs
- Isoefficiency func (due to concurrency) is optimal, i.e. $\Theta(p)$ only if $C(W) = \Theta(W)$
if $C(W) < \Theta(W)$, isoefficiency fn (due to concurrency) $> \Theta(p)$

SOURCES OF PARALLEL OVHD T_0 :

- ① Interprocessor communication
- ② Load imbalance → cannot predict / need to sync / seq. components
- ③ Extra computation → fastest seq. algo may be too hard to parallelize
 ↳ (serial:) reuse results, (parallel:) cannot reuse, as generated by different processors \therefore extra computation e.g. FFT

- performance of scaled down algo. may be different for different assignments of (virtual \rightarrow real) processors

eg $[n \times n][]$ on p -proc HC $\left\{ \begin{array}{l} \rightarrow p \text{ square blocks} \\ \rightarrow p \text{ stripes of } n/p \text{ rows} \end{array} \right.$

ISOEFFICIENCY METRIC OF SCALABILITY

- useful to determine rate at which problem size must increase w.r.t. $\#p$ to keep the efficiency fixed
- Problem size \equiv $\#$ computation steps in best seq. algo on 1 proc.

Overhead $T_0(W, p) = pT_p - W$

$$T_0 = pT_p - T_s$$

- Problem size: regardless of problem, doubling size \Rightarrow 2 times the computation.

For matrix mult, $W = \theta(n^3)$

matrix addn, $W = \theta(n^2)$

$$\therefore W = T_s$$

• Isoefficiency fn:

- in 1 expr, captures characteristics of \parallel algo & \parallel architecture
- test performance of \parallel pgm on a few processors;] not viable
predict performance on larger $\#$ processors
- characterizes amt of parallelism inherent in a parallel algo
- study behavior of \parallel system w.r.t. changes in HW parameters eg speed of CPUs, channels
- useful for \parallel algos for which we cannot derive a value of \parallel runtime.