

HPC – Unit 4: Performance Metrics & Matrix Computation

May–June 2023 (Paper [6004]-493)

Q3 a) Parallel Matrix-Matrix Multiplication Algorithm with Example [7 Marks]

Matrix-matrix multiplication $C = A \times B$ where A is an $m \times k$ matrix and B is a $k \times n$ matrix. The sequential algorithm requires $O(m \cdot k \cdot n)$ operations. Parallel algorithms distribute either rows, columns, or sub-blocks across processors.

Method 1: Row-wise Distribution

Each of the p processors is assigned m/p rows of matrix A . Every processor needs all of B to compute its rows of C .

- Phase 1 (Broadcast B): One-to-all broadcast B to all processors — cost: $O(\log p \times k \cdot n)$.
- Phase 2 (Local multiply): Each processor computes (m/p) rows of C independently.
- Parallel time: $T_p \approx (m \cdot k \cdot n)/p + \log p \times k \cdot n$.

Example: Multiply two 4×4 matrices on 2 processors.

$A = \begin{bmatrix} 1 & 2 & 3 & 4 \\ 5 & 6 & 7 & 8 \\ 9 & 0 & 1 & 2 \\ 3 & 4 & 5 & 6 \end{bmatrix}$	$B = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$
--	--

P_0 gets rows 0,1 of $A \rightarrow$ computes rows 0,1 of $C =$ rows 0,1 of A (since $B=I$)

P_1 gets rows 2,3 of $A \rightarrow$ computes rows 2,3 of C

Method 2: Block (Cannon's) Algorithm

For square $n \times n$ matrices on p processors arranged in a $\sqrt{p} \times \sqrt{p}$ mesh, Cannon's algorithm avoids replicating the full matrix by using circular shifts.

- Sub-matrices are distributed: processor (i,j) gets block $A[i,j]$ and $B[i,j]$.
- Initial alignment: shift rows of A left by i positions and columns of B up by j positions.
- Each step: processors multiply their local sub-blocks and accumulate into C . Then A shifts left by 1 and B shifts up by 1.
- After \sqrt{p} steps, each processor has its complete $C[i,j]$ block.

Parallel time (Cannon's): $T_p = O(n^3/p) + O(\sqrt{p} \times n^2/p \times t_{\text{comm}})$

For large n , communication is dominated by computation — good scalability.

Note: For SPPU exams, explaining the row-distribution approach with a small numerical example (2×2 or 4×4 with 2 or 4 processors) is usually sufficient for full marks.

Q3 b) Performance Metrics for Parallel Systems

[6 Marks]

Several metrics are used to evaluate how well a parallel program utilizes available processors:

1. Speedup (S)

$$S = T_s / T_p$$

where T_s = best sequential execution time, T_p = parallel execution time on p processors.

- Ideal speedup = p (linear). In practice, $S < p$ due to overhead.
- $S > p$ (superlinear speedup) can occur due to cache effects — each processor's working set fits in cache.

2. Efficiency (E)

$$E = S / p = T_s / (p \times T_p)$$

Ranges from 0 to 1; ideal = 1 (100% efficient).

- Efficiency drops as p increases because overhead (communication, synchronisation, idle time) grows.

3. Cost (C)

$$C = p \times T_p$$

Total processor-time product. A cost-optimal algorithm has $C = O(T_s)$.

- Cost optimality means the parallel algorithm uses total work proportional to the sequential algorithm — no wasted processor cycles.

4. Speedup Factor (Amdahl's Perspective)

$$S \leq 1 / (f + (1-f)/p)$$

where f = fraction of program that is inherently sequential (cannot be parallelised).

5. Utilisation (U)

$$U = (\text{Total useful work}) / (p \times T_p) = E$$

- Utilisation equals efficiency — measures what fraction of processor time is spent on useful computation (vs. overhead).

6. Execution Rate

- Measured in FLOPS (floating-point operations per second) or MIPS.
- Parallel FLOPS = $p \times (\text{sequential FLOPS per processor}) \times \text{efficiency}$.

7. Redundancy

$R = \text{Total operations in parallel} / \text{Operations in best sequential algorithm}$

- $R \geq 1$ always; $R = 1$ means no redundant work. Some parallel algorithms intentionally do extra work to reduce communication.

Q3 c) Minimum Execution Time and Minimum Cost-Optimal Execution Time

[4 Marks]

Minimum Execution Time

The minimum execution time T_{\min} is the fastest possible wall-clock time achievable using any number of processors, regardless of cost. It is obtained when we use as many processors as the problem's parallelism allows.

$$T_{\min} = \lim_{p \rightarrow \infty} T_p$$

For a problem with a sequential fraction f (Amdahl's Law):

$$T_{\min} = f \times T_s \quad (\text{the inherently serial portion cannot be parallelised})$$

- Using more processors beyond a point gives diminishing returns because the sequential bottleneck dominates.

Minimum Cost-Optimal Execution Time

A parallel algorithm is cost-optimal if the total cost $C = p \times T_p$ equals the sequential time T_s (asymptotically). The minimum cost-optimal execution time T_{\min_cost} is the fastest time achievable while keeping cost optimal:

T_{\min_cost} is achieved when p is chosen such that:
 $p \times T_p = \Theta(T_s)$ (cost matches sequential work)

For example, sorting n elements sequentially: $T_s = O(n \log n)$

Parallel bitonic sort on n processors: $T_p = O(\log^2 n)$

Cost = $n \times \log^2 n > n \log n \rightarrow$ not cost-optimal.

Using $p = n/\log n$ processors gives $T_p = O(\log n \times \log(n/\log n)) \approx O(\log^2 n)$ with cost $O(n \log n) \rightarrow$ cost-optimal.

- Cost-optimal algorithms are preferred in practice because they are economically efficient — you get maximum speedup per unit of hardware cost.

Note: Minimum execution time is about speed; minimum cost-optimal execution time balances speed and resource cost. Both are key concepts in algorithmic complexity for parallel systems.

Q4 a) Granularity and Its Effects on Performance

[7 Marks]

Granularity refers to the ratio of computation to communication in a parallel program. It determines how work is partitioned among processors.

Types of Granularity

- Fine-grained: Very small tasks per processor. Many small computations, frequent communication. High communication overhead relative to computation.
- Coarse-grained: Large tasks per processor. Few, infrequent communications. Reduced parallelism (fewer concurrent pieces).
- Medium-grained: A balance that gives good speedup and manageable communication overhead — often the sweet spot.

Granularity = (Computation time per processor) / (Communication time)

Fine: ratio $\ll 1$ (communication dominates — bad)

Coarse: ratio $\gg 1$ (underutilised parallelism — also potentially bad)

Optimal: ratio $\approx O(1 \text{ to } 10)$ depending on interconnect speed

Effect of Granularity on Performance — Adding n Numbers on p Processors

Consider summing n numbers distributed across p processors:

Phase 1 (local sum): Each processor sums its n/p numbers.

$$T_{\text{local}} = n/p \times t_{\text{add}}$$

Phase 2 (tree reduction): $\log_2(p)$ rounds of communicate-and-add

$$T_{\text{comm}} = \log_2(p) \times (t_s + t_w)$$

Total: $T_p = n/p \times t_{\text{add}} + \log_2(p) \times t_s$

Speedup: $S = (n \times t_{\text{add}}) / (n/p \times t_{\text{add}} + \log_2(p) \times t_s)$

- As n increases (coarser grain for fixed p): speedup approaches p — good efficiency.
- As p increases for fixed n (finer grain): local work shrinks, communication dominates, efficiency drops.
- As $n/p \rightarrow 0$: $T_p \approx \log_2(p) \times t_s$ — all communication, no useful work.

Note: This example directly shows why granularity matters: a fine-grained decomposition (small n , large p) can actually run slower than a coarser partition. Always tune granularity based on the ratio of computation to communication bandwidth.

Q4 b) Sources of Overhead in Parallel Systems

[6 Marks]

Overhead T_o is the extra time spent by a parallel program over and above the best sequential execution. Formally, $T_o = p \times T_p - T_s$. It comes from several sources:

1. Inter-Process Communication

Moving data between processors is the dominant overhead in most parallel programs. Each message incurs a startup cost (latency) and a per-byte transfer cost. As p increases, more messages are needed, driving up this overhead.

2. Load Imbalance

If the work is not evenly distributed, some processors sit idle while others are still computing. The parallel time is determined by the slowest (busiest) processor. Load imbalance overhead = $(T_{\max} - T_{\text{avg}}) \times p$.

3. Synchronisation Overhead

Barriers, locks, and other synchronisation primitives cause processors to wait for each other. Even a short delay at a barrier ripples across all p processors.

4. Extra Computation

Some parallel algorithms do more total work than the best sequential algorithm (e.g., sorting networks). This redundant computation is overhead. Quantified by the redundancy metric R .

5. Task Scheduling and Management

In dynamic parallel frameworks (OpenMP task pools, work-stealing schedulers), there is overhead in creating, scheduling, and terminating tasks. For very fine-grained tasks this can dominate.

6. Memory Access and Cache Effects

In shared-memory systems, false sharing (two processors share a cache line but access different variables) causes unnecessary cache invalidations, stalling processors.

Summary: $T_o = T_{\text{comm}} + T_{\text{idle}} + T_{\text{sync}} + T_{\text{redundant}} + T_{\text{sched}}$

Total parallel time: $T_p = (T_s + T_o) / p$

Efficiency: $E = T_s / (T_s + T_o) \rightarrow \text{lower overhead} \Rightarrow \text{higher efficiency}$

Q4 c) Scaling Down (Downsizing) a Parallel System

[4 Marks]

Scaling down refers to taking a parallel algorithm designed for p processors and running it on $p' < p$ processors (fewer processors than originally designed for). This is common when the full machine is not available or when running smaller problem sizes.

Technique: Virtual Processors (Folding)

Each physical processor simulates multiple virtual processors. If the original algorithm used $p = 64$ processors and only $p' = 8$ are available, each physical processor handles $p/p' = 8$ virtual nodes.

Example: 1D prefix-sum designed for $p=8$ nodes, now on $p'=2$ nodes.

Each physical node simulates 4 virtual nodes.

Physical node 0 handles virtual nodes {0,1,2,3}
 Physical node 1 handles virtual nodes {4,5,6,7}

Intra-node communication becomes local computation (no network needed).

Inter-node communication only when virtual nodes on different physical nodes must talk.

Effect on performance:

- Computation time increases by factor p/p' (each processor does more work).
- Communication time may decrease (some messages become local).
- Net effect: $T_{p'} \leq (p/p') \times T_p$ — often better than expected because of eliminated communication.
- Efficiency can actually improve when scaling down, since fewer boundary communications are needed.

Note: Scaling down is an important practical technique for running HPC codes on smaller clusters without rewriting the algorithm. It demonstrates the ISO-efficiency relationship: maintaining efficiency as both p and problem size change.

May–June 2024 (Paper [6263]-94)

Q3 a) Sources of Overhead in Parallel Systems

[7 Marks]

[REPEATED] – See Q4 b) in May–June 2023 above for the complete answer.

Q3 b) Effect of Granularity on Performance – Adding n Numbers on p Processors

[6 Marks]

[REPEATED] – See Q4 a) in May–June 2023 above for the complete answer (includes the n -number addition example).

Q3 c) Amdahl's Law and Gustafson's Law

[4 Marks]

Amdahl's Law

Amdahl's Law gives an upper bound on the speedup achievable by parallelising a program, given that a fraction f of the program is inherently sequential:

$$S(p) \leq 1 / (f + (1-f)/p)$$

$$\text{As } p \rightarrow \infty: S_{\max} = 1/f$$

Example: If $f = 0.1$ (10% sequential), $S_{\max} = 10x$ — even with infinite processors.

- Implication: even a small sequential section sharply limits speedup.
- Pessimistic view — assumes problem size is fixed as p grows.

Example table for $f = 0.1$:

Processors (p)	Speedup S
1	1.0

2	1.82
4	3.08
8	4.71
16	6.40
∞	10.0

Gustafson's Law (Scaled Speedup)

Gustafson argued that as we add more processors, we typically also increase the problem size (scaled workload). The sequential fraction then becomes a smaller portion of the total work:

$$\text{Scaled speedup } S_{\text{scaled}} = p - f(p - 1)$$

where f = fraction of parallel runtime that is sequential.

Equivalently: $S_{\text{scaled}} = p - \alpha(p-1)$ where α = sequential fraction of parallel runtime.

- Gustafson's view is optimistic: with larger problems, the parallel portion dominates and speedup can scale linearly with p .
- More realistic for HPC workloads where scientists always want to solve bigger problems when given more processors.

Comparison:

Aspect	Amdahl's Law	Gustafson's Law
Problem size	Fixed	Scales with p
Focus	Strong scaling	Weak scaling
Speedup bound	$1/f$ (hard ceiling)	$p - f(p-1)$ (grows with p)
Outlook	Pessimistic	Optimistic
When accurate	Real-time constraints	Scientific computing

Note: In SPPU exams, always write both formulas, give a numerical example for each, and compare them. Examiners expect both laws in a single question.

Q4 a) Performance Metrics for Parallel Systems

[7 Marks]

[REPEATED] – See Q3 b) in May–June 2023 above for the complete answer. All metrics (Speedup, Efficiency, Cost, Utilisation, Execution Rate, Redundancy) are covered there.

Q4 b) Parallel Matrix-Matrix Multiplication Algorithm with Example

[6 Marks]

[REPEATED] – See Q3 a) in May–June 2023 above for the complete answer.

Q4 c) Scalability of Parallel Systems

[4 Marks]

Scalability measures how effectively a parallel system (algorithm + architecture) maintains its efficiency

as both the problem size n and the number of processors p grow simultaneously.

ISO-Efficiency Function

A scalable parallel system keeps efficiency $E = E_0$ (constant) as p increases, provided the problem size W is scaled up at the same rate. The ISO-efficiency function $W(p)$ gives the minimum problem size needed to maintain efficiency E_0 on p processors:

$$E = T_s / (T_s + T_o) = E_0 \\ \Rightarrow T_o = ((1 - E_0)/E_0) \times W$$

$$\text{So: } W \geq (E_0 / (1 - E_0)) \times T_o(p)$$

Example: If $T_o = O(p \log p)$ (overhead for a broadcast-heavy algorithm), then W must grow as $O(p \log p)$ to keep efficiency constant. This means the ISO-efficiency function is $W(p) = O(p \log p)$ — scalable.

Types of Scalability

- Strongly scalable: Efficiency stays constant even for fixed problem size. Requires T_o to be small relative to T_s .
- Weakly scalable: Efficiency stays constant only if problem size grows with p . Most practical parallel algorithms fall here.
- Not scalable: T_o grows faster than W for any reasonable scaling of W . Overhead overwhelms useful work.

Factors Affecting Scalability

- Communication overhead — grows with p , degrades scalability.
- Load imbalance — unequal work distribution worsens with heterogeneous tasks.
- Synchronisation points — global barriers become bottlenecks as p grows.
- Memory per processor — if each processor needs access to global data, memory pressure can limit p .

Note: A good rule of thumb: an algorithm with $T_o = O(p)$ has a linear ISO-efficiency function $W(p) = O(p)$ — it scales well. An algorithm with $T_o = O(p^2)$ is poor; W must grow quadratically to maintain efficiency.

Additional Concepts & Quick Reference

Key Formulas Summary

Metric	Formula	Ideal Value
Speedup	$S = T_s / T_p$	p (linear)
Efficiency	$E = S/p = T_s/(p \cdot T_p)$	1 (100%)
Cost	$C = p \times T_p$	$O(T_s)$ (cost-optimal)
Overhead	$T_o = p \cdot T_p - T_s$	0
Amdahl bound	$S \leq 1/(f + (1-f)/p)$	$1/f$ as $p \rightarrow \infty$
Gustafson	$S = p - f(p-1)$	Grows with p

Matrix Multiplication – Complexity Summary

- Sequential: $O(n^3)$ for $n \times n$ matrices (naive); $O(n^{2.37})$ with Strassen-like algorithms.
- Parallel row-distribution on p processors: $T_p = O(n^3/p) + O(n^2 \log p)$ — good for large n .
- Cannon's algorithm on $\sqrt{p} \times \sqrt{p}$ mesh: $T_p = O(n^3/p) + O(n^2/\sqrt{p})$ — near-linear speedup.
- Cost-optimal when $p \leq n^2 / \log n$ (so that communication doesn't dominate).

Nov–Dec 2025 (Paper [6584]-81)

Q3 a) Parallel Matrix-Matrix Multiplication Algorithm with Example [7 Marks]

[REPEATED] – Complete answer with row-distribution method, Cannon's algorithm, worked 4x4 example, and parallel time formulas is in: Unit 4 Answer Doc → May–June 2023 → Q3 a).

Q3 b) Circular Shift Operation on Mesh and Hypercube Networks [6 Marks]

A circular shift is a collective communication where each node i sends its data to node $(i + q) \bmod p$ and receives from node $(i - q + p) \bmod p$. On a mesh or hypercube the efficient implementation exploits the topology's multi-dimensional structure.

Circular Shift on a Mesh

A $p_r \times p_c$ mesh has nodes identified by (row, col). A general 2D circular shift of (q_r, q_c) — meaning node (i, j) sends to node $((i+q_r) \bmod p_r, (j+q_c) \bmod p_c)$ — is handled in two sequential phases, exactly as described in the Unit 3 note for Cannon's algorithm.

Phase 1 (Row shift): Every row performs a simultaneous 1D circular shift by q_c positions. All rows act independently and in parallel, so this takes q_c communication steps along the horizontal links.

Phase 2 (Column shift): Every column performs a simultaneous 1D circular shift by q_r positions. All columns act independently and in parallel, taking q_r steps along the vertical links.

Example: 3×3 mesh, shift $(q_r=1, q_c=1)$

Before shift:

$(0,0)=A$ $(0,1)=B$ $(0,2)=C$
 $(1,0)=D$ $(1,1)=E$ $(1,2)=F$
 $(2,0)=G$ $(2,1)=H$ $(2,2)=I$

After row shift $q_c=1$ (each element moves one column right, wraps):

$(0,0)=C$ $(0,1)=A$ $(0,2)=B$
 $(1,0)=F$ $(1,1)=D$ $(1,2)=E$
 $(2,0)=I$ $(2,1)=G$ $(2,2)=H$

After column shift $q_r=1$ (each element moves one row down, wraps):

$(0,0)=I$ $(0,1)=G$ $(0,2)=H$
 $(1,0)=C$ $(1,1)=A$ $(1,2)=B$
 $(2,0)=F$ $(2,1)=D$ $(2,2)=E$

Node (i,j) now holds what was originally at $((i-1+3)\%3, (j-1+3)\%3)$ ✓

Cost: $T_{\text{mesh}} = (q_r + q_c)(t_s + m \cdot t_w)$

For Cannon's algorithm $q_r = q_c = 1$ each round: $T = 2(t_s + m \cdot t_w)$ per step.

Circular Shift on a Hypercube

A d -dimensional hypercube with $p = 2^d$ nodes identifies each node by its d -bit binary address. A circular shift by q positions maps node i to node $(i + q) \bmod p$. The beauty of the hypercube is that any two nodes differ in exactly those bits where their XOR is non-zero, meaning a shift of q steps can be implemented using at most $d = \log_2 p$ communication steps — one per differing bit.

The key insight is that the shift displacement q can be expressed in binary, and each set bit in q corresponds to one communication step across that hypercube dimension:

Algorithm: Hypercube circular shift by q

```
decompose q into binary:  $q = b_{d-1} \dots b_1 b_0$ 
for each bit  $k$  where  $b_k = 1$ :
    each node sends its current data to the neighbour along dimension  $k$ 
    (i.e.,  $\text{partner} = \text{current\_id} \oplus 2^k$ )
    simultaneously receive from the same partner
```

Example: 8-node hypercube ($d=3$), shift $q=5$ (binary 101)

Bit 0 set: send across dimension 0 ($\text{partner} = \text{id} \oplus 1$)

Bit 2 set: send across dimension 2 ($\text{partner} = \text{id} \oplus 4$)

Total: 2 steps for shift of 5 positions.

General cost: $T_{\text{hypercube}} = \text{popcount}(q) \times (t_s + m \cdot t_w)$

$\text{popcount}(q)$ = number of set bits in $q \leq \log_2(p)$

Worst case ($q = p/2 = 2^{d-1}$): only 1 step!

Average case: $\approx d/2$ steps

Comparison: Mesh vs Hypercube for Circular Shift

Aspect	Mesh ($p_r \times p_c$)	Hypercube (2^d)
Steps for shift (q_r, q_c)	$q_r + q_c$ steps	$\text{popcount}(q) \leq \log_2 p$ steps
Best case	1 step (axis-aligned)	1 step (power-of-2 shift)
Worst case	$p_r + p_c - 2$ steps	$\log_2 p$ steps
Physical links per node	2 or 4 (left/right/up/down)	$\log_2 p$ (one per dimension)
Application	Cannon's matrix multiply	Bitonic sort, FFT butterfly

Note: Cannon's matrix multiplication is the canonical application of circular shift on a mesh — the initial alignment shifts A left by i positions and B up by j positions. The hypercube shift is used in FFT butterfly communications and in sorting networks like bitonic sort.

Q3 c) How to Improve Speed of Communication Operations

[4 Marks]

[REPEATED] – Full answer covering pipeline communication, recursive doubling, spanning trees, ring AllReduce, and hardware offloading is in: Unit 3 Answer Doc → May–June 2024 → Q2 c). Note: this question appears in Q3 of the Nov 2025 paper but covers Unit 3 syllabus content.

Q4 a) Granularity and Effects on Performance

[7 Marks]

[REPEATED] – Full answer with definitions, fine/coarse/medium granularity, and the n -number addition worked example is in: Unit 4 Answer Doc → May–June 2023 → Q4 a).

Q4 b) Sources of Overhead in Parallel Programs**[6 Marks]**

[REPEATED] – Full answer covering communication overhead, load imbalance, synchronisation, extra computation, task scheduling, and false sharing is in: Unit 4 Answer Doc → May–June 2023 → Q4 b).

Q4 c) Performance Metrics of Parallel Systems**[4 Marks]**

[REPEATED] – Full answer with all metrics (Speedup, Efficiency, Cost, Utilisation, Execution Rate, Redundancy) and their formulas is in: Unit 4 Answer Doc → May–June 2023 → Q3 b).

May–June 2025 (Paper [6404]-94)**Q3 a) Performance Metrics of Parallel Systems****[6 Marks]**

[REPEATED] – See Unit 4 Answer Doc → May–June 2023 → Q3 b) for all metrics with formulas.

Q3 b) Matrix-Matrix Multiplication in Details**[6 Marks]**

[REPEATED] – Full answer with row-distribution and Cannon's algorithm is in: Unit 4 Answer Doc → May–June 2023 → Q3 a).

Q3 c) Minimum and Cost-Optimal Execution Time**[5 Marks]**

[REPEATED] – Full answer with formulas and sorting example is in: Unit 4 Answer Doc → May–June 2023 → Q3 c).

Q4 a) Matrix-Vector Multiplication: Row-wise 1D, 2D Partitioning, and Comparison**[9 Marks]**

Matrix-vector multiplication computes $y = A \cdot x$, where A is an $n \times n$ matrix and x is an n -element vector. Since every element y_i requires one full dot product (n multiplications and additions), there are n^2 operations total, taking $O(n^2)$ sequentially. The key challenge in parallelisation is that every row of A needs the full vector x , which creates a communication bottleneck.

i) Row-wise 1D Partitioning

The simplest approach is to partition the rows of A across p processors. Processor k receives rows $(k \cdot n/p)$ through $((k+1) \cdot n/p - 1)$ of A , along with the corresponding rows of y to compute. The critical issue is that every processor needs all n elements of x to compute its assigned rows of y .

Setup: $n \times n$ matrix A , n -vector x , computing $y = A \cdot x$ on p processors.
 Processor k owns: rows $[k \cdot n/p .. (k+1) \cdot n/p - 1]$ of A , same rows of y .

Step 1 — Distribute A : scatter rows of A to appropriate processors.
 Each processor receives n/p rows $\times n$ elements = n^2/p elements.
 Cost: $O(n^2/p)$ to scatter.

Step 2 — Broadcast x : every processor needs the full vector x .
 One-to-all broadcast of n elements to all p processors.
 Cost: $O(n \log p)$ for a hypercube, $O(n \times p)$ for a ring.

Step 3 — Local computation: each processor computes its n/p dot products.
 $T_{\text{comp}} = (n/p) \times 2n = 2n^2/p$ multiplications/additions.

Total parallel time:

$$T_p = O(n^2/p) + O(n \log p) \quad [\text{scatter} + \text{broadcast} + \text{computation}]$$

Speedup: $S \approx p$ for large n (computation dominates when $n \gg p \cdot \log p$)

Limitation: the broadcast of x (step 2) sends the entire n -element vector to every processor. For large n this is acceptable, but as p grows the broadcast overhead can dominate.

ii) 2D Partitioning

In 2D partitioning, both the matrix A and the vector x are partitioned. The p processors are arranged in a $\sqrt{p} \times \sqrt{p}$ logical grid. Processor (i, j) owns the $n/\sqrt{p} \times n/\sqrt{p}$ sub-block $A[i,j]$ of the matrix and a sub-segment of x .

Setup: $\sqrt{p} \times \sqrt{p}$ processor grid. Processor (i,j) owns:

Sub-block of A : rows $[i \cdot n/\sqrt{p} .. (i+1) \cdot n/\sqrt{p} - 1]$,

cols $[j \cdot n/\sqrt{p} .. (j+1) \cdot n/\sqrt{p} - 1]$

A segment of x : elements $[j \cdot n/\sqrt{p} .. (j+1) \cdot n/\sqrt{p} - 1]$

Step 1 — Column broadcast of x : each column of processors broadcasts

its segment of x to all processors in the same column.

Each processor now has the n/\sqrt{p} elements of x it needs for its columns of A .

$$\text{Cost: } O((n/\sqrt{p}) \log \sqrt{p}) = O((n/\sqrt{p}) \times \log(p)/2)$$

Step 2 — Local computation: processor (i,j) computes a partial dot product

for each of its n/\sqrt{p} assigned rows of y :

$$y_{\text{partial}}[\text{row}] = \sum A[\text{row}, \text{col}] \times x[\text{col}] \quad \text{for the } n/\sqrt{p} \text{ columns it owns.}$$

$$T_{\text{comp}} = O((n/\sqrt{p})^2) = O(n^2/p) \text{ per processor.}$$

Step 3 — Row reduction: processors in the same row sum their partial results to produce the final y values for their row range.

$$\text{Cost: } O((n/\sqrt{p}) \times \log \sqrt{p}) = O((n/\sqrt{p}) \times \log(p)/2)$$

Total parallel time:

$$T_p = O(n^2/p) + O((n/\sqrt{p}) \times \log p)$$

iii) Comparison of 1D and 2D Partitioning

Aspect	Row-wise 1D	2D Partitioning
Matrix distribution	n/p rows per processor	$n/\sqrt{p} \times n/\sqrt{p}$ block per processor
Communication	Broadcast full x : $O(n \log p)$	Broadcast x segment + reduce: $O((n/\sqrt{p}) \log p)$
Computation	$O(n^2/p)$ per processor	$O(n^2/p)$ per processor
Memory per processor	$O(n^2/p)$ for A , $O(n)$ for full x	$O(n^2/p)$ for A , $O(n/\sqrt{p})$ for x segment
Scalability	Communication grows as $O(n \log p)$	Communication grows as $O((n/\sqrt{p}) \log p)$
Preferred when	p is small, simple to implement	p is large, memory-constrained
ISO-efficiency function	$W(p) = O(p \log^2 p)$	$W(p) = O(p \log^2 p / \sqrt{p})$ — better

The fundamental reason 2D partitioning wins at large p is that it reduces the communication volume per processor from $O(n)$ (full x) to $O(n/\sqrt{p})$ (a fraction of x). For a $\sqrt{p} \times \sqrt{p}$ grid, each column only needs to share its n/\sqrt{p} segment of x within its column, not broadcast n elements globally. The reduction in communication volume is a factor of \sqrt{p} , which becomes significant as p grows.

Note: When $n \gg \sqrt{p}$, both methods give comparable speedup. The 2D method becomes critically important when running thousands of processors on a modern supercomputer, where $O(n)$ broadcasts would overwhelm the network. This is why modern libraries like ScaLAPACK and Intel MKL use 2D block-cyclic layouts for distributed dense linear algebra.

Q4 b) Dense Matrix Algorithms: Matrix-Vector Multiplication and Matrix-Matrix Multiplication [8 Marks]

i) Dense Matrix-Vector Multiplication (MVM)

MVM is the building block of iterative solvers (conjugate gradient, GMRES) and neural network inference. The complete algorithm for parallel MVM was covered in Q4 a) above. Here we summarise the key algorithmic points:

- The fundamental operation is $y_i = \sum_{j=0}^{n-1} A_{ij} \times x_j$ for each row i .
- Data access pattern: each row of A is read once, but the entire x vector must be accessed for every row. This makes x the communication bottleneck.
- Arithmetic intensity: $O(2n)$ operations per $O(n)$ data read per row — relatively low arithmetic intensity, making MVM memory-bandwidth bound rather than compute-bound on modern hardware.
- Parallel efficiency: For 1D partitioning, efficiency $E = T_s / (p \times T_p) = n^2 / (n^2 + pn \cdot \log p) \rightarrow 1$ as $n \rightarrow \infty$ for fixed p .

Parallel MVM summary:

Sequential time: $T_s = 2n^2$

1D parallel time: $T_p = 2n^2/p + n \cdot \log_2(p) \cdot t_{s/m}$ (broadcast overhead)

2D parallel time: $T_p = 2n^2/p + (n/\sqrt{p}) \cdot \log_2(p) \cdot t_{s/m}$ (better)

ii) Dense Matrix-Matrix Multiplication (MMM)

[Core algorithm REPEATED] – Full MMM answer with row-distribution and Cannon's algorithm is in: Unit 4 Answer Doc → May–June 2023 → Q3 a). The following extends that with the dense algorithm framing:

For dense MMM $C = A \times B$ (all $n \times n$), the three standard parallel approaches are:

Approach 1 — Simple Row Distribution: Distribute rows of A across p processors; broadcast entire B to all. Each processor computes its n/p rows of C locally.

Cost: broadcast $B = O(n^2 \log p) + \text{local compute} = O(n^3/p)$

Total $T_p = O(n^3/p + n^2 \log p)$

Efficient when $n \gg p$, i.e., $n \gg p \cdot \log p$

Approach 2 — Cannon's Algorithm (2D mesh): The n^2/p sub-block approach where sub-matrices are aligned by circular shifts, then each processor performs \sqrt{p} local multiplications of $n/\sqrt{p} \times n/\sqrt{p}$ blocks.

Cost: circular shifts $= O(\sqrt{p} \times n^2/p) + \text{local compute} = O(n^3/p)$

Total $T_p = O(n^3/p + n^2/\sqrt{p} \times t_w)$

Key advantage: no full matrix broadcast; total communication $= O(n^2/\sqrt{p})$

Approach 3 — SUMMA (Scalable Universal Matrix Multiply): The industry-standard algorithm used in ScaLAPACK. Like Cannon's but uses broadcasts along rows and columns instead of shifts, making it more flexible for non-square processor grids.

SUMMA: each of n/b outer-product steps broadcasts one column panel of A and one row panel of B to the corresponding processor row/column.

Block size b is tuned to the network bandwidth.

Cost: $O(n^3/p + n^2 \log p / \sqrt{p})$

Algorithm	Communication	Computation	Best For
Row distribution	$O(n^2 \log p)$ — broadcast full B	$O(n^3/p)$	Small p , simple code
Cannon's (mesh)	$O(n^2/\sqrt{p} \times t_w)$ — shifts only	$O(n^3/p)$	Square mesh, HPC clusters
SUMMA	$O(n^2 \log p / \sqrt{p})$	$O(n^3/p)$	Production HPC, ScaLAPACK

Note: The term 'dense' distinguishes algorithms that treat all matrix entries as non-zero (no sparsity exploited). Dense algorithms are communication-bound for large p — the n^3/p computation becomes so fast on modern FPUs that inter-processor communication is the bottleneck. This is why minimising communication volume (Cannon's, SUMMA) is the primary goal of dense linear algebra on parallel computers.

Additional Concepts & Quick Reference

1D vs 2D Partitioning: When to Use Which

The choice between 1D and 2D partitioning hinges on the ratio of n to p . For 1D partitioning, each processor needs $O(n)$ extra memory for the full copy of x , independent of p . For 2D partitioning, each processor only needs $O(n/\sqrt{p})$ extra memory for its x segment. As p grows (more processors, smaller problem per processor), 1D partitioning becomes increasingly wasteful of both memory and communication bandwidth. The crossover point is roughly when $p > \sqrt{n}$ — at that level, 2D partitioning is definitively better.

Matrix Algorithm Complexity Summary

Problem	Sequential	Parallel (1D)	Parallel (2D/Cannon)
Matrix-vector multiply	$O(n^2)$	$O(n^2/p + n \log p)$	$O(n^2/p + (n/\sqrt{p}) \log p)$
Matrix-matrix multiply	$O(n^3)$	$O(n^3/p + n^2 \log p)$	$O(n^3/p + n^2/\sqrt{p})$

$O(n^3/p + n^2/\sqrt{p})$