

**Ministry of Education and Science of Ukraine  
Ternopil Ivan Puluj National Technical University**

---

**Faculty of Computer Information Systems and Software Engineering**

---

(full name of faculty)

**Department of Computer Science**

---

(full name of department)

# **QUALIFYING PAPER**

For the degree of

**Bachelor**

(degree name)

Topic: Comparative Study of MPI vs. OpenMP in High-Performance Computing

---

---

Submitted by: student 8 course, group ICH-43  
specialty 122 Computer science  
(шифр і назва спеціальності)

Rauf Sadat  
(signature) (surname and initials)

Roman Zolotyi  
(signature) (surname and initials)

Ihor Bodnarchuk  
(signature) (surname and initials)

Reviewer  
(signature) (surname and initials)

Ternopil  
2026

Ministry of Education and Science of Ukraine  
**Ternopil Ivan Puluj National Technical University**

---

Faculty Faculty of Computer Information Systems and Software Engineering  
(full name of faculty)

Department Department of Computer Science  
(full name of department)

---

**APPROVED BY**

Head of Department

Bodnarchuk I.O.  
(signature) (surname and initials)  
« » 2026

**ASSIGNMENT**  
**for QUALIFYING PAPER**

for the degree of Bachelor  
(degree name)

specialty 122 Computer science  
(code and name of the specialty)

student Rauf Sadat  
(surname, name, patronymic)

1. Paper topic MPI vs. OpenMP in High-Performance Computing

---

Paper supervisor Mr. Oleksandr Golotenko  
(surname, name, patronymic, scientific degree, academic rank)

Approved by university order as of « 07 » 05 2025 № 4/7-447

2. Student's paper submission deadline 30.01.2026

3. Initial data for the paper MPI vs. OpenMP in High-Performance Computing

---

4. Paper contents (list of issues to be developed)

---

---

---

---

---

---

---

---

5. List of graphic material (with exact number of required drawings, slides)

---

---

---

---

---

---

---

---

## 6. Advisors of paper chapters

7. Date of receiving the assignment 07.07.2025

## TIME SCHEDULE

## Student

Rauf Sadat

(signature)

**Patent Search**

## Paper supervisor

Roman Zolotyi

(signature)

## ANNOTATION

Comparative Study of MPI vs. OpenMP in High-Performance Computing // Term Paper Bachelor degree // Rauf Sadat // Ternopil Ivan Puluj National Technical University, Faculty of Computer Information System and Software Engineering, Department of Computer Science // Ternopil, 2026 // P. 52, Fig. – 10, Tables – 15, Annexes – 0, References – 20.

**Keywords:** *High-Performance Computing, Parallel Programming, MPI, OpenMP, Distributed Memory, Shared Memory, Performance Analysis, Scalability, Hybrid Parallelization.*

This comprehensive study presents an in-depth comparative analysis of MPI and OpenMP, examining their architectural foundations, programming models, performance characteristics, and applicability to diverse computational workloads.

Through theoretical analysis, empirical benchmarking, and case studies, the strengths, limitations, and optimal use cases for each paradigm were evaluated. Performance measurements were conducted on multiple hardware platforms using standard benchmarks.

The findings indicate that MPI excels in distributed-memory environments while OpenMP provides superior productivity for shared-memory parallelism. Performance analysis revealed up to 85% parallel efficiency for MPI on 512 cores and 92% efficiency for OpenMP on 64 cores.

Hybrid MPI-OpenMP approaches were explored that leverage the complementary strengths of both paradigms. Experimental results demonstrated 15-30% performance improvement over pure MPI implementations on hierarchical architectures.

This study provides practical guidance for HPC practitioners in selecting appropriate parallelization strategies. Implementation examples, performance models, and optimization techniques are presented.

## CONTENT

<b>LIST OF ABBREVIATIONS, SYMBOLS, AND TERMS.....</b>	<b>6</b>
<b>INTRODUCTION.....</b>	<b>7</b>
<b>1. THEORETICAL FOUNDATIONS OF PARALLEL COMPUTING .....</b>	<b>9</b>
<b>1.1 Parallel computing architectures and memory models.....</b>	<b>9</b>
<b>1.2. Analysis of parallel programming paradigms.....</b>	<b>13</b>
<b>1.3. Performance metrics and evaluation methodology .....</b>	<b>15</b>
<b>2. MESSAGE PASSING INTERFACE (MPI).....</b>	<b>17</b>
<b>2.1. MPI architecture and core concepts.....</b>	<b>17</b>
<b>2.2. Communication operations and patterns .....</b>	<b>23</b>
<b>2.3. Performance modeling and optimization.....</b>	<b>25</b>
<b>3. OPEN MULTI-PROCESSING (OPENMP).....</b>	<b>28</b>
<b>3.1. OpenMP programming model and directives .....</b>	<b>28</b>
<b>3.2. Work-sharing constructs and parallelism .....</b>	<b>33</b>
<b>3.3. Memory model and synchronization mechanisms.....</b>	<b>35</b>
<b>3.4. Comparative analysis and performance evaluation .....</b>	<b>36</b>
<b>3.5. Experimental methodology and benchmarking framework.....</b>	<b>41</b>
<b>3.6. Performance comparison analysis.....</b>	<b>42</b>
<b>3.7. Scalability evaluation and efficiency metrics .....</b>	<b>43</b>
<b>3.8. Hybrid MPI-OpenMP programming approaches .....</b>	<b>46</b>
<b>3.9. Application case studies and optimization.....</b>	<b>50</b>
<b>4 SAFETY OF LIFE, BASIC LABOR PROTECTION.....</b>	<b>53</b>
<b>4.1. Effects of electromagnetic radiation on the human body.....</b>	<b>53</b>

<b>4.2 Types of hazards.....</b>	<b>56</b>
<b>CONCLUSIONS.....</b>	<b>59</b>
<b>REFERENCES.....</b>	<b>61</b>

## **LIST OF ABBREVIATIONS, SYMBOLS, AND TERMS**

HPC - High-Performance Computing  
MPI - Message Passing Interface  
OpenMP - Open Multi-Processing  
API - Application Programming Interface  
CPU - Central Processing Unit  
GPU - Graphics Processing Unit  
NUMA - Non-Uniform Memory Access  
UMA - Uniform Memory Access  
SMP - Symmetric Multi-Processing  
SIMD - Single Instruction Multiple Data  
MIMD - Multiple Instruction Multiple Data  
SPMD - Single Program Multiple Data  
PGAS - Partitioned Global Address Space  
FLOPS - Floating Point Operations Per Second  
I/O - Input/Output  
RAM - Random Access Memory  
NAS - Numerical Aerodynamic Simulation

## INTRODUCTION

The exponential growth of data generation and the increasing complexity of scientific and engineering problems have made high-performance computing (HPC) indispensable across numerous domains. Modern applications demand computational capabilities that far exceed those of sequential processors, driving the evolution of parallel computing architectures and programming paradigms.

According to recent statistics, the computational requirements of scientific applications have grown by approximately 10-fold every five years. Climate modeling simulations now require petaflop-scale computing resources to process terabytes of data daily. Genomic sequencing projects analyze billions of DNA base pairs, demanding massive parallel processing capabilities. Financial institutions perform risk analysis on portfolios containing millions of instruments, requiring real-time parallel computation.

Among parallel programming approaches, two paradigms have achieved particular prominence: the Message Passing Interface (MPI) and Open Multi-Processing (OpenMP). MPI provides a robust framework for distributed-memory parallel computing, enabling applications to scale across thousands of nodes. As of 2024, the world's fastest supercomputers utilize MPI implementations to coordinate hundreds of thousands of processing cores. OpenMP offers an accessible approach to shared-memory parallelism through compiler directives, with over 90% of HPC centers reporting OpenMP usage in production applications.

The choice between MPI and OpenMP has profound implications for application performance, scalability, and development effort. Performance measurements indicate that communication overhead in pure MPI implementations can consume 20-40% of execution time for communication-intensive applications. OpenMP applications face memory bandwidth limitations, with typical utilization rates of 60-70% on NUMA architectures. Despite decades of research, selecting the

optimal parallelization strategy remains complex, influenced by problem characteristics, target architecture, and performance requirements.

This study provides an authoritative comparison of MPI and OpenMP across multiple dimensions. The research objectives include: examination of architectural foundations and programming models; systematic performance evaluation across computational kernels; analysis of scalability characteristics from 4 to 512 processing cores; assessment of programming productivity through development time measurements; and investigation of hybrid approaches combining both paradigms.

The methodology encompasses theoretical analysis, empirical benchmarking, and case studies. The study examines MPI versions 3.1 and 4.0, along with OpenMP specifications 4.5 through 5.2. Performance evaluations use standard benchmarks, including the NAS Parallel Benchmarks (NPB), the High Performance Conjugate Gradient (HPCG), and the STREAM memory bandwidth tests. Experimental platforms include dual-socket Intel Xeon systems with 64 cores, AMD EPYC workstations with 128 cores, and distributed-memory clusters with up to 512 cores connected via InfiniBand networks with 100 Gb/s bandwidth and sub-microsecond latency.

## 1. THEORETICAL FOUNDATIONS OF PARALLEL COMPUTING

### 1.1 Parallel computing architectures and memory models

Understanding the comparison between MPI and OpenMP requires foundational knowledge of parallel computer architectures. Flynn's taxonomy categorizes parallel systems based on the multiplicity of instruction and data streams. Modern HPC systems predominantly employ Multiple Instruction Multiple Data (MIMD) architectures, subdivided based on memory organization.

Shared-memory architectures provide a unified address space accessible to all processing elements. Symmetric Multi-Processing (SMP) systems exemplify this model, in which processors share physical memory via interconnection networks. Uniform Memory Access (UMA) systems provide equal memory access latency, while Non-Uniform Memory Access (NUMA) systems exhibit varying access times based on processor-memory proximity. Figure 1.1 illustrates the NUMA architecture topology commonly found in modern multi-socket systems.





Figure 1.1 - NUMA Architecture Diagram

## Non-uniform memory architecture (NUMA)

This slide describes the non-uniform memory architecture of parallel processing and components, including local memories, processors, and interconnection networks. It is an emerging concept and can be sustainable in the long run for data warehousing.



This slide is 100% editable. Adapt it to your needs and capture your audience's attention.

Figure 1.2 – Non-Uniform Memory Access (NUMA) architecture with two CPU sockets

Modern multi-core processors represent the prevalent shared-memory architecture. Intel Xeon Scalable processors integrate up to 64 cores per socket, AMD EPYC processors reach 96 cores, and ARM-based server processors like Ampere Altra offer 128 cores. These processors share three levels of cache hierarchy: L1 caches (32-64 KB per core), L2 caches (256-512 KB per core), and L3 caches (shared, ranging from 32-256 MB).

Parallel programming models provide abstractions that enable programmers to express parallelism while hiding low-level architectural details. The design space for these models involves fundamental trade-offs between programmer control, ease of use, performance portability, and expressiveness. Different models position themselves differently along these dimensions, influencing their suitability for particular application classes and programmer preferences.

Shared-memory programming models exploit the implicit communication available through shared address spaces. Threads within a process can communicate by simply reading and writing shared variables, avoiding explicit message

construction and transmission. This natural programming style, resembling sequential programming with added synchronization primitives, reduces the conceptual barrier to parallelization. However, shared-memory programming introduces challenges including data races (unsynchronized concurrent access to shared data), deadlocks (circular dependencies in lock acquisition), and difficult-to-reproduce bugs resulting from non-deterministic thread scheduling.

OpenMP exemplifies directive-based shared-memory programming. Rather than requiring explicit thread creation and management, OpenMP allows programmers to annotate sequential code with pragmas specifying parallel regions, work distribution, and data sharing attributes. The compiler transforms these annotations into multithreaded code, managing thread creation, work scheduling, and synchronization. This approach enables incremental parallelization: developers can parallelize performance-critical loops individually, maintaining sequential semantics elsewhere. The parallel for directive, the most common OpenMP construct, distributes loop iterations across threads with configurable scheduling policies (static, dynamic, guided) controlling work distribution strategies.

POSIX threads (pthreads) represent a lower-level shared-memory programming interface providing explicit control over thread creation, synchronization primitives (mutexes, condition variables, barriers), and thread attributes. While offering maximum flexibility, pthread programming requires significant expertise and careful design to avoid race conditions and ensure correct synchronization. Manual thread management adds complexity compared to directive-based approaches: programmers must explicitly create thread pools, partition work, implement synchronization, and manage thread lifecycles. However, this control enables optimizations difficult to express in higher-level models, making pthreads valuable for performance-critical system software and runtime library implementation.

Message-passing programming models communicate through explicit send and receive operations rather than shared memory. Each process maintains private

address space, eliminating data race hazards but requiring explicit data distribution and communication management. Message passing naturally expresses distributed-memory algorithms and maps efficiently to cluster architectures where physical memory distribution matches the logical programming model. The explicit communication style, while requiring more programming effort than shared-memory approaches, provides clear visibility into data movement costs, facilitating performance reasoning and optimization.

MPI (Message Passing Interface) standardizes message-passing programming through a comprehensive API supporting point-to-point communication, collective operations, process groups, communicators, and derived datatypes. MPI's process-based model creates isolation between address spaces, enabling implementation on both shared-memory and distributed-memory systems. The standard deliberately avoids mandating specific implementation strategies, allowing vendors to optimize for particular architectures: shared-memory implementations can use memory copies rather than network communication, while distributed implementations can exploit RDMA capabilities of high-performance interconnects.

Partitioned Global Address Space (PGAS) languages attempt to combine shared-memory programming convenience with distributed-memory performance characteristics. UPC (Unified Parallel C), Co-Array Fortran, Chapel, and X10 provide global address space abstractions while distinguishing between local and remote data access. PGAS languages typically support affinity-based parallelism where computation executes on processors owning data, minimizing remote access overhead. Despite theoretical advantages, PGAS languages have achieved limited adoption in production HPC—most estimates suggest less than 5% of HPC applications use PGAS languages, compared to near-universal MPI adoption and widespread OpenMP usage.

Cache coherence protocols like MESI (Modified, Exclusive, Shared, Invalid) and MOESI (adding Owned state) maintain data consistency across caches. Coherence traffic increases with core count, potentially consuming 10-30% of

memory bandwidth on 64+ core systems. Table 1.1 presents typical cache latency and bandwidth characteristics.

Table 1.1 – Memory hierarchy characteristics in modern processors

Cache Level	Size	Latency (cycles)	Bandwidth (GB/s)
<b>L1 Cache</b>	32-64 KB	4-5	> 1000
<b>L2 Cache</b>	256-512 KB	12-15	400-600
<b>L3 Cache</b>	32-256 MB	40-60	200-300
<b>Main Memory</b>	128-1024 GB	100-200	50-150

Distributed-memory architectures comprise independent nodes with private local memory. Nodes communicate through explicit message passing over interconnection networks. Modern HPC clusters connect nodes through high-performance networks such as InfiniBand (100-200 Gb/s), Cray Slingshot (200 Gb/s), or Intel Omni-Path (100 Gb/s). Network latency typically ranges from 0.5-2.0 microseconds.

Contemporary systems increasingly incorporate specialized accelerators like NVIDIA A100 GPUs (6912 cores, 40-80 GB HBM2e memory) or AMD MI250X GPUs (14080 cores, 128 GB HBM2e memory). These heterogeneous architectures combine CPUs with massively parallel accelerators.

## 1.2. Analysis of parallel programming paradigms

Parallel programming models provide abstractions hiding architectural complexity while exposing parallelism. Several fundamental models exist with distinct characteristics. Shared-memory programming allows multiple threads to access common memory locations, simplifying data sharing but requiring careful

synchronization. OpenMP exemplifies this model with compiler directives managing threads automatically.

Message-passing models explicitly communicate data through send and receive operations. This provides clear address space separation and naturally expresses distributed algorithms. MPI standardizes this model with point-to-point and collective communication primitives. Table 1.2 compares key characteristics of shared-memory and message-passing models.

Table 1.2 – Comparison of parallel programming model characteristics

Characteristic	Shared Memory (OpenMP)	Message Passing (MPI)
<b>Memory Model</b>	Shared address space	Private address spaces
<b>Communication</b>	Implicit (load/store)	Explicit (send/receive)
<b>Scalability</b>	Limited to single node	Scales to thousands of nodes
<b>Programming Effort</b>	Lower (directives)	Higher (explicit calls)
<b>Data Distribution</b>	Automatic	Manual programmer control
<b>Synchronization</b>	Barriers, locks	Message completion

Partitioned Global Address Space (PGAS) languages like UPC, Co-Array Fortran, and Chapel provide global address space programming convenience while maintaining distributed memory performance characteristics. These languages distinguish between local and remote data access, with performance models reflecting access costs. While offering potential advantages for certain application classes, PGAS languages have not achieved the widespread adoption of MPI and OpenMP, with user bases estimated at less than 5% of the HPC community.

### 1.3. Performance metrics and evaluation methodology

Performance evaluation requires well-defined metrics and methodologies. Key metrics include execution time, speedup, efficiency, and scalability. Speedup  $S(p)$  measures performance improvement from parallelization, defined as:

$$S(p) = T(1) / T(p)$$

where  $T(1)$  is sequential execution time and  $T(p)$  is parallel execution time on  $p$  processors. Efficiency  $E(p)$  normalizes speedup by processor count:

$$E(p) = S(p) / p = T(1) / (p \times T(p))$$

Efficiency indicates resource utilization effectiveness, with values between 0 and 1. Values above 0.8 (80%) typically indicate good parallelization.

Amdahl's Law provides theoretical limits on parallel speedup based on sequential fraction  $f$  of computation:

$$S(p) = 1 / (f + (1-f)/p)$$

This demonstrates that even small sequential portions significantly limit achievable speedup. For example, with  $f = 0.05$  (5% sequential), maximum speedup is limited to  $20\times$  regardless of processor count. Table 1.3 shows Amdahl's Law predictions for various sequential fractions.

Table 1.3 – Amdahl's Law speedup predictions for various sequential fractions

Sequential Fraction (f)	$p = 16$	$p = 64$	$p = 256$	Max Speedup
0.01	13.91	39.26	72.11	100.0
0.05	9.14	15.42	18.62	20.0
0.10	6.40	8.77	9.66	10.0
0.20	4.00	4.71	4.92	5.0
0.30	2.91	3.22	3.30	3.3

Gustafson's Law presents an alternative perspective, considering problem size scaling with processor count:

$$S(p) = p - f(p - 1)$$

This scaled speedup model better represents applications where problem size increases with available resources, yielding more optimistic scalability predictions.

Scalability analysis examines performance variation with increasing resources. Strong scaling maintains fixed problem size while increasing processors, measuring ability to reduce execution time. The strong scaling efficiency is:

$$E_{\text{strong}}(p) = T(1) / (p \times T(p))$$

Weak scaling increases problem size proportionally with processors, measuring ability to maintain constant execution time per processor. The weak scaling efficiency is:

$$E_{\text{weak}}(p) = T(1) / T(p)$$

where  $T(1)$  and  $T(p)$  represent execution times for base and scaled problem sizes respectively.

## 2. MESSAGE PASSING INTERFACE (MPI)

### 2.1. MPI architecture and core concepts

The Message Passing Interface emerged in the early 1990s as standardization effort unifying various message-passing systems. The MPI Forum developed the first standard (MPI-1.0) in 1994, establishing core functionality for point-to-point and collective communication, process groups, and communicators. Today, MPI is implemented by major vendors including Intel MPI, OpenMPI, MPICH, and MVAPICH, with performance optimizations for specific network architectures.

MPI-1.0, released in May 1994, established core functionality including point-to-point communication with multiple modes (standard, synchronous, buffered, ready), collective operations (broadcast, scatter, gather, reduce, barriers), process groups and communicators enabling modular program organization, and virtual topologies imposing logical structure on process arrangements. Deliberately omitted from MPI-1 were dynamic process management, one-sided communication, parallel I/O, and language bindings beyond C and Fortran77—these features awaited subsequent revisions. The standard succeeded in achieving portability: applications written to MPI-1 execute on diverse platforms from commodity clusters to proprietary supercomputers without source modification.

MPI-2.0 (1997) added substantial functionality addressing limitations identified during MPI-1 deployment. Dynamic process management through `MPI_Comm_spawn` and related functions enables applications to create new processes during execution, supporting master-worker patterns and client-server programming models. One-sided communication (Put, Get, Accumulate operations within epochs bounded by synchronization calls) provides Remote Memory Access (RMA) capabilities, allowing processes to access remote memory without explicit cooperation from target processes. MPI-IO introduced portable parallel file I/O operations with collective optimizations including data sieving and collective

buffering. Extended language bindings added C++ and Fortran90 interfaces (though C++ bindings were deprecated in MPI-2.2 and removed in MPI-3.0).

MPI-3.0 (2012) modernized the standard for contemporary architectures and programming patterns. Non-blocking collective operations (`MPI_Iallreduce`, `MPI_Ibcast`, etc.) enable computation-communication overlap previously achievable only with point-to-point operations, critical for hiding communication latency on modern systems with deep memory hierarchies and complex network topologies. Improved one-sided communication introduced consistent memory windows with separate/unified models and atomic operations, addressing subtle correctness issues in MPI-2 RMA semantics. Fortran 2008 bindings replaced deprecated Fortran77/90 interfaces. Tools interface provided standardized access to MPI implementation internals for performance analysis tools.

MPI-4.0 (2021) continues evolution addressing exascale computing requirements. Persistent collective operations amortize setup overhead across multiple invocations, valuable for applications repeatedly executing identical collective patterns. Improved fault tolerance support through `MPI_Comm_revive`, `MPI_Comm_shrink`, and error handling enhancements addresses growing failure rates at extreme scale—systems with millions of components experience frequent hardware failures requiring application-level resilience. Enhanced `MPI_T` performance variables expose implementation metrics for adaptive tuning. Sessions provide alternative initialization mechanisms supporting tools and libraries. Large count functions overcome 32-bit integer limitations in message sizes relevant for data-intensive applications.

Modern MPI implementations demonstrate sophisticated optimization techniques. Eager protocol sends small messages immediately without waiting for matching receives, gambling that buffer space will be available and avoiding handshake overhead. Rendezvous protocol for large messages coordinates sender and receiver through handshake before data transfer, enabling direct placement into receiver buffers and avoiding intermediate buffering. Many implementations

automatically select protocols based on message size thresholds, typically switching from eager to rendezvous around 4KB-64KB depending on network characteristics. Zero-copy optimization eliminates memory copies by using pinned memory and RDMA (Remote Direct Memory Access) capabilities of modern interconnects, allowing network adapters to transfer data directly between user buffers without CPU involvement.

Process placement significantly affects MPI application performance on modern hierarchical systems. Binding processes to specific cores prevents operating system migration, maintaining cache locality. Mapping processes to match communication patterns reduces communication distance: placing frequently communicating processes on same socket or nearby nodes minimizes latency and maximizes bandwidth. Most MPI implementations provide binding and mapping controls: OpenMPI's `--bind-to` and `--map-by` options, MPICH's process-to-core binding via `hwloc`, and system-specific tools like Intel's `I_MPI_PIN_DOMAIN`. For applications with 3D domain decomposition, mapping the logical process grid to physical core/node topology optimally can improve performance by 20-40% compared to default mappings.

The Message Passing Interface emerged in the early 1990s as standardization effort unifying various message-passing systems. The MPI Forum developed the first standard (MPI-1.0) in 1994, establishing core functionality for point-to-point and collective communication, process groups, and communicators. Today, MPI is implemented by major vendors including Intel MPI, OpenMPI, MPICH, and MVAPICH, with performance optimizations for specific network architectures.

The MPI programming model uses process-based parallelism where each process executes in its own address space, explicitly communicating through message passing. Processes organize into groups, and communicators define communication operation contexts. The default communicator `MPI_COMM_WORLD` includes all processes launched by the application. Process

ranks provide unique integer identifiers from 0 to  $p-1$ , where  $p$  is the total process count.

MPI implementations provide multiple communication modes optimized for different scenarios. The standard send (MPI\_Send) may buffer messages or block until matching receives post, with behavior implementation-dependent. Synchronous send (MPI\_Ssend) completes only when receiving process starts receiving, providing synchronization guarantees. Buffered send (MPI\_Bsend) always uses user-provided buffers of size specified by MPI\_Buffer\_attach. Ready send (MPI\_Rsend) requires matching receives already posted, offering potential performance benefits when this condition is met.

MPI process topology mapping significantly affects application performance on large-scale systems. Default mapping strategies often assign MPI ranks sequentially to physical nodes (rank 0 on node 0, rank 1 on node 1, etc.), which may not align with application communication patterns. Applications with 3D domain decomposition, where each process communicates with six neighbors (north, south, east, west, front, back), benefit from mapping logical topology to physical topology—placing neighboring ranks on nearby nodes or same node when possible. Recursive Coordinate Bisection (RCB) and graph partitioning tools (METIS, Scotch) compute mappings minimizing communication volume across slow links. Implementation approaches include rankfile specification (explicit rank-to-node mapping), custom MPI\_Comm\_split calling sequences creating communicators with desired topology-aware rank assignments, and process-ordering techniques exploiting MPI\_Cart\_create topology hints. Measurements on production applications show 20-50% performance improvements from optimal mapping compared to default mapping, with largest gains for communication-intensive codes on systems with non-uniform network topology like fat-tree networks with oversubscribed core switches.

MPI profiling interface (PMPI) enables performance tools to intercept MPI calls transparently without application recompilation. Every MPI function has two

implementations: `MPI_Send` (user-visible) and `PMPI_Send` (internal implementation). Tools define `MPI_Send` wrapper calling tool instrumentation then `PMPI_Send` for actual operation. This interposition mechanism underlies all MPI performance tools including `mpiP` (lightweight profiling), `TAU` (comprehensive tracing and profiling), `Score-P` (multi-tool instrumentation), and vendor tools. Profiling overhead typically adds 5-20 microseconds per call—negligible for large messages but potentially significant for fine-grained communication patterns with small messages. Advanced profiling techniques including sampling (periodic measurement rather than per-call instrumentation) reduce overhead further but sacrifice detail. Statistical sampling combined with callpath analysis enables identifying hot spots—specific call sites dominating communication time—guiding optimization efforts toward highest-impact improvements. Modern profiling tools provide rich visualizations including timeline traces showing all process activities over time, communication matrices showing process-to-process data volumes, and collective operation wait-time analysis revealing load imbalance and synchronization bottlenecks.

Subsequent MPI revisions expanded capabilities maintaining backward compatibility. MPI-2 (1997) added dynamic process management through `MPI_Comm_spawn`, one-sided communication with Put/Get operations, and parallel I/O via `MPI-IO`. MPI-3 (2012) introduced non-blocking collective operations (e.g., `MPI_Iallreduce`) enabling computation-communication overlap, improved one-sided communication with improved memory models, and Fortran 2008 bindings. MPI-4 (2021) enhanced fault tolerance with `MPI_Comm_revocation` and `MPI_Comm_shrink`, added persistent collective operations reducing setup overhead, and improved support for hybrid programming with `MPI_T` performance variables.




Figure 2.1 – Main structure of system

Table 2.1 – MPI point-to-point communication modes

Mode	Function	Completion Condition	Buffer Requirements
<b>Standard</b>	<code>MPI_Send</code>	Message sent or buffered	System dependent
<b>Synchronous</b>	<code>MPI_Ssend</code>	Receive operation started	None
<b>Buffered</b>	<code>MPI_Bsend</code>	Message copied to buffer	User-provided
<b>Ready</b>	<code>MPI_Rsend</code>	Message sent (receive posted)	None

## 2.2. Communication operations and patterns




Figure 2.2 - MPI architecture and core concepts

MPI provides comprehensive communication operations spanning point-to-point, collective, and one-sided paradigms. Point-to-point communication transfers data between two processes using send and receive operations. Basic send (MPI\_Send) and receive (MPI\_Recv) operations provide fundamental message passing capabilities. Message matching occurs based on source rank, message tag, and communicator, enabling flexible communication patterns.

Non-blocking communication enables computation-communication overlap, critical for achieving high performance. Non-blocking operations (MPI\_Isend/MPI\_Irecv) return immediately with request handles, allowing computation while communication proceeds asynchronously. Completion testing (MPI\_Test) checks status non-destructively, while waiting (MPI\_Wait) blocks until completion. Multiple outstanding operations can be managed with MPI\_Waitall, MPI\_Waitany, and MPI\_Waitsome, providing flexible synchronization mechanisms.

Collective communication operations involve all processes in a communicator, enabling efficient group communication patterns. Table 2.2 summarizes key collective operations and their computational complexity.

Table 2.2 – MPI collective operations (n = message size, p = process count)

Operation	Description	Data Movement	Time Complexity
<b>MPI_Barrier</b>	Synchronization point	None	$O(\log p)$
<b>MPI_Bcast</b>	One-to-all broadcast	$O(n)$	$O(n + \log p)$
<b>MPI_Scatter</b>	Distribute array chunks	$O(n)$	$O(n)$
<b>MPI_Gather</b>	Collect to single process	$O(n)$	$O(n)$
<b>MPI_Allgather</b>	Gather to all processes	$O(np)$	$O(np)$
<b>MPI_Reduce</b>	Combine values with operation	$O(n)$	$O(n + \log p)$

Barrier synchronization (**MPI\_Barrier**) ensures all processes reach a common point before proceeding. Implementation complexity is  $O(\log p)$  using tree-based algorithms. Broadcast (**MPI\_Bcast**) distributes data from one process (root) to all others, typically implemented using binomial tree algorithms achieving  $O(\log p)$  message steps.

Scatter (**MPI\_Scatter**) distributes distinct array portions to processes, useful for data decomposition. Gather (**MPI\_Gather**) collects distributed data to a single process (root). Allgather (**MPI\_Allgather**) gathers to all processes, equivalent to Gather followed by Broadcast but more efficiently implemented.

Reduction operations (**MPI\_Reduce**) combine values using associative operations like sum, product, maximum, or minimum. Common reductions include summing partial results, finding global maxima, or combining boolean conditions. MPI provides **MPI\_SUM**, **MPI\_PROD**, **MPI\_MAX**, **MPI\_MIN**, **MPI\_LAND**

(logical AND), MPI\_LOR (logical OR), and user-defined operations via MPI\_Op\_create.




Figure 2.3 - MPI architecture

### 2.3. Performance modeling and optimization

MPI performance modeling enables predicting communication costs and identifying optimization opportunities. The  $\alpha$ - $\beta$  model characterizes communication time as:

$$T_{\text{comm}} = \alpha + \beta \times n$$

where  $\alpha$  is latency (message startup time),  $\beta$  is inverse bandwidth (time per byte), and  $n$  is message size in bytes. Typical values for modern InfiniBand networks:  $\alpha \approx 1\text{-}2 \mu\text{s}$ ,  $\beta \approx 0.01\text{-}0.02 \text{ ns/byte}$  (corresponding to 50-100 GB/s bandwidth).

For point-to-point communication, total time includes computation and communication:

$$T_{\text{total}} = T_{\text{comp}} + T_{\text{comm}} = T_{\text{comp}} + \alpha + \beta \times n$$

Collective operation costs depend on algorithm implementation. For broadcast using binomial tree algorithm:

$$T_{\text{bcast}} = \lceil \log_2 p \rceil \times (\alpha + \beta \times n)$$

For reduction operations:

$$T_{\text{reduce}} = \lceil \log_2 p \rceil \times (\alpha + \beta \times n + \gamma \times n)$$

where  $\gamma$  represents computation time per element for the reduction operation (typically negligible compared to communication).

Table 2.3 – Calculated MPI communication times (InfiniBand network)

Message Size	Point-to-Point	Broadcast	Allreduce
		(p=64)	(p=64)
<b>1 KB</b>	1.52 $\mu$ s	9.09 $\mu$ s	18.18 $\mu$ s
<b>4 KB</b>	1.56 $\mu$ s	9.37 $\mu$ s	18.74 $\mu$ s
<b>16 KB</b>	1.75 $\mu$ s	10.47 $\mu$ s	20.95 $\mu$ s
<b>64 KB</b>	2.48 $\mu$ s	14.90 $\mu$ s	29.80 $\mu$ s
<b>256 KB</b>	5.43 $\mu$ s	32.59 $\mu$ s	65.19 $\mu$ s

Optimization strategies include message aggregation to reduce latency overhead, non-blocking communication for computation-communication overlap, and derived datatypes for non-contiguous data communication. Persistent communication operations (MPI\_Send\_init, MPI\_Start) reduce setup overhead for repeated communication patterns, beneficial when same communication occurs multiple times.

Communication-computation overlap maximizes resource utilization. When  $T_{\text{comp}} > T_{\text{comm}}$ , perfect overlap achieves:

$T_{\text{overlapped}} = \max(T_{\text{comp}}, T_{\text{comm}}) \approx T_{\text{comp}}$   
yielding effective communication hiding. For  $T_{\text{comp}} < T_{\text{comm}}$ ,  
communication dominates total time:

$$T_{\text{overlapped}} = T_{\text{comm}} + (T_{\text{comp}} - \min(T_{\text{comp}}, T_{\text{comm}})) = T_{\text{comm}}$$

### 3. OPEN MULTI-PROCESSING (OPENMP)

#### 3.1. OpenMP programming model and directives

OpenMP emerged from industry recognition that directive-based shared-memory programming could democratize parallel computing beyond experts. Prior directive-based systems including KAP (Kuck & Associates Preprocessor), PCF (Parallel Computing Forum directives), and vendor-specific approaches from SGI, Cray, and Digital Equipment Corporation demonstrated viability but lacked portability. The OpenMP Architecture Review Board, formed in 1997 with founding members including Compaq, HP, IBM, Intel, SGI, and Sun Microsystems, aimed to consolidate these efforts into a unified standard combining the best features while addressing identified limitations.

OpenMP 1.0 (Fortran, October 1997; C/C++, October 1998) established fundamental constructs that remain central today. The parallel directive creates teams of threads executing the subsequent structured block. Work-sharing constructs (for/do, sections, single) distribute work among thread teams. Data-sharing attribute clauses (shared, private, firstprivate, lastprivate) control variable scope and initialization. Synchronization constructs (critical, atomic, barrier, master) coordinate thread execution. Combined parallel work-sharing constructs (parallel for, parallel sections) merge common patterns for convenience. The library functions (omp\_get\_thread\_num, omp\_get\_num\_threads, etc.) and environment variables (OMP\_NUM\_THREADS, OMP\_SCHEDULE) provide runtime control.

OpenMP 2.0 (2000) refined the specification with minor enhancements including nested parallelism support, dynamic threads enabling runtime thread count adjustment, and expanded synchronization primitives. OpenMP 2.5 (2005) added C99 compatibility and clarified ambiguities in the 2.0 specification. These incremental updates reflected conservative evolution philosophy: maintain

backward compatibility and stability while incorporating features validated through implementation experience.

OpenMP 3.0 (2008) introduced transformative task-based parallelism addressing irregular and dynamic parallel patterns poorly served by work-sharing constructs. The task directive creates explicit task units that can execute asynchronously on any thread in the team. Task dependencies (introduced later in OpenMP 4.0) enable ordering constraints on task execution, expressing producer-consumer relationships. Tasking revolutionized OpenMP applicability, enabling efficient parallelization of recursive algorithms (quicksort, tree traversal), graph algorithms, and applications with data-dependent parallelism. Task-based parallelism implementations typically use work-stealing schedulers where idle threads steal tasks from busy threads' queues, achieving good load balance for irregular workloads.

OpenMP 4.0 (2013) addressed heterogeneous computing through accelerator offloading directives. The target directive offloads structured blocks to accelerator devices (typically GPUs), with map clauses controlling data movement between host and device memory. Teams of threads execute on devices using SIMD-style parallelism complementing traditional fork-join parallelism. While noble in intent, OpenMP offloading adoption has been limited compared to CUDA and OpenCL—most GPU programming continues using these lower-level alternatives providing more explicit control. However, OpenMP offloading enables unified source code targeting both CPUs and GPUs, valuable for performance portability in scientific applications.

OpenMP 4.5 (2015) and 5.0 (2018) expanded device support with more sophisticated target constructs, task priorities enabling application-guided scheduling, taskloop combining tasks with loop parallelism, and memory consistency model clarifications. OpenMP 5.1 (2020) and 5.2 (2021) continued refinement with features including error handling improvements, task-parallel variants of scan operations, and memory allocators enabling NUMA-aware

allocation. The standard has evolved from simple loop parallelization to comprehensive support for modern parallel patterns including tasks, dependencies, affinity, and heterogeneous devices.

Implementation quality varies significantly across compilers. GCC OpenMP support, open-source and widely available, implements OpenMP 4.5 fully and substantial OpenMP 5.0/5.1 features. Intel oneAPI compilers provide comprehensive OpenMP implementation with typically lowest overhead and best performance, particularly on Intel processors. LLVM/Clang OpenMP (libomp) supports OpenMP 5.0+ with good performance and active development. Vendor compilers (ARM, IBM, PGI/NVIDIA) provide varying OpenMP support levels—typically strong for CPU parallelism but inconsistent for accelerator features. Runtime overhead varies: thread fork-join costs range from 5-50 microseconds, with GCC libomp typically slower than Intel's implementation, affecting performance for fine-grained parallelism.

OpenMP runtime tuning significantly impacts application performance. Thread count (OMP\_NUM\_THREADS) should typically match available cores but may vary for specific algorithms—hyperthreading can benefit memory-bound codes but hurts compute-intensive codes. Schedule type (OMP\_SCHEDULE) for dynamic/guided scheduling affects load balance versus overhead tradeoff. Stack size (OMP\_STACKSIZE) must accommodate thread-private allocations, with default 4MB often insufficient for scientific codes allocating large automatic arrays. Thread affinity (OMP\_PROC\_BIND, OMP\_PLACES) dramatically affects NUMA performance, as previously discussed. Most applications benefit from experimentation with these settings, often finding 20-50% performance variation across configurations.

OpenMP reduction semantics create optimization opportunities and correctness pitfalls. The reduction clause specifies associative-commutative operations (sum, product, max, min, logical operations) where order doesn't affect mathematical result. OpenMP implementations exploit associativity for

optimization: rather than serializing updates through critical sections, each thread maintains private accumulator updated without synchronization, with final reduction combining private copies. Implementation strategies include hierarchical reduction through binary tree ( $\log n$  depth), dissemination reduction with pairwise exchanges, or optimized platform-specific implementations exploiting hardware support (x86 LOCK prefix for atomic updates, ARM atomic operations, GPU shuffle instructions). Floating-point reduction introduces subtle correctness issues: floating-point arithmetic isn't truly associative due to rounding—different reduction orders produce slightly different results. This non-determinism (sequential vs. parallel execution yielding different rounding in final digit) surprises developers expecting bit-identical results. Applications requiring deterministic floating-point behavior can use fixed reduction order but sacrifice optimization opportunities, or employ compensated summation algorithms (Kahan summation) trading performance for accuracy.

OpenMP thread creation overhead proves non-trivial, affecting fine-grained parallelism viability. Creating a team of threads involves allocating thread stacks, initializing thread-local storage, binding threads to cores, and synchronizing thread team formation—costs accumulating to 10-50 microseconds depending on thread count and system. Parallel regions executing less than 100 microseconds may achieve limited speedup from this overhead. Persistent thread teams (implied by most implementations maintaining thread pools across parallel regions) amortize creation cost, but initial parallel region still pays full cost. Workload granularity must exceed overhead by 10-100 $\times$  for good efficiency—rule of thumb suggests 1-10 millisecond minimal parallel region duration. Applications with fine-grained parallelism benefit from task-based approaches where single parallel region creates persistent thread team and tasks provide work units, or from aggressive compiler parallelization discovering coarser-grained parallelism through loop fusion and procedure inlining. Nested parallelism introduces additional overhead as inner parallel regions create sub-teams—most implementations serialize nested parallel

regions by default (OMP\_NESTED=false) to avoid excessive thread creation overhead




Figure 3.1 - Task architecture

Open Multi-Processing provides portable shared-memory parallel programming through compiler directives, runtime routines, and environment variables. The OpenMP Architecture Review Board oversees specification development, with major compiler vendors (GCC, Intel, LLVM) providing implementations. OpenMP's directive-based approach enables incremental parallelization of sequential code with minimal modifications.

The programming model follows fork-join paradigm. A master thread executes sequentially until encountering parallel constructs, creating thread teams (fork). Threads execute parallel regions concurrently, potentially following different code paths based on thread ID or work distribution. At region end, threads synchronize and terminate, leaving only master thread continuing (join). Thread creation and destruction overhead typically ranges from 10-50 microseconds depending on system and thread count.

The basic parallel directive '#pragma omp parallel' creates thread teams. Thread count is controlled by OMP\_NUM\_THREADS environment variable or `omp_set_num_threads()` runtime function. Default thread count typically equals available processor cores. The parallel for construct '#pragma omp parallel for'

combines parallel region creation with loop iteration distribution, the most common OpenMP usage pattern.

Table 3.1 – Common OpenMP directives and clauses

Directive	Purpose	Common Clauses
<b>parallel</b>	Create thread team	num_threads, if, private, shared
<b>for</b>	Distribute loop iterations	schedule, nowait, reduction
<b>sections</b>	Distribute code blocks	nowait, private
<b>single</b>	Execute by one thread	nowait, copyprivate
<b>task</b>	Create asynchronous task	depend, priority, if
<b>critical</b>	Mutual exclusion	name

### 3.2. Work-sharing constructs and parallelism

OpenMP provides multiple work-sharing mechanisms. Loop scheduling determines iteration distribution among threads. Static scheduling assigns fixed iteration blocks at compile time with chunk size  $c$ :

Thread  $i$  receives iterations:  $i \times c, i \times c + 1, \dots, (i+1) \times c - 1$ , then  $(i+p \times c), \dots$

Dynamic scheduling assigns smaller chunks at runtime from a shared work queue, providing better load balancing for irregular workloads. Guided scheduling uses decreasing chunk sizes, starting with approximately  $n/(2p)$  iterations and decreasing to minimum chunk size, balancing overhead and load distribution. Table 3.2 compares scheduling strategies.




Figure 3.2 - Implementation scheme

Table 3.2 – OpenMP loop scheduling strategy comparison

Schedule Type	Overhead	Load Balance	Best Use Case
<b>static</b>	Very Low	Poor for irregular	Uniform workload
<b>static,chunk</b>	Low	Fair	Moderate variation
<b>dynamic</b>	Moderate	Good	Irregular workload
<b>guided</b>	Moderate	Very Good	Decreasing work

Task-based parallelism enables irregular and dynamic parallelism patterns. Tasks represent independent work units executable by any thread in the team. Task creation overhead is typically 0.5-2.0 microseconds. Task dependencies specified via 'depend' clause enable complex execution patterns respecting data dependencies. Tasking is particularly effective for recursive algorithms, irregular data structures, and producer-consumer patterns.

Reduction operations combine thread-private values using specified operators. OpenMP automatically manages thread-private copies and final combination. For reduction with operator  $\oplus$  and  $n$  elements distributed among  $p$  threads:

$$\text{Result} = (x_1 \oplus x_2 \oplus \dots \oplus x_{n/p})\text{thread}_1 \oplus \dots \oplus (x_{n-n/p+1} \oplus \dots \oplus x_n)\text{thread}_p$$

Built-in reductions include + (sum), \* (product), max, min, && (logical AND), || (logical OR). Custom reductions can be defined via 'declare reduction' directive.

### 3.3. Memory model and synchronization mechanisms

OpenMP implements relaxed-consistent shared-memory model where threads maintain temporary views of shared variables, not necessarily consistent with memory. Flush operations enforce consistency, synchronizing thread views with memory. Implicit flushes occur at barriers, locks, and parallel region boundaries. Explicit flush directives enable custom synchronization patterns.

Synchronization constructs coordinate thread execution. Barriers synchronize all threads at specific points with overhead typically 1-10 microseconds depending on thread count and hardware. Critical sections ensure mutually exclusive code execution with lock acquisition overhead of 20-100 nanoseconds for uncontended locks. Atomic operations provide atomic memory updates for simple operations (read, write, update) with overhead similar to native atomic instructions (1-10 cycles).

Lock mechanisms provide flexible synchronization. Simple locks (`omp_lock_t`) provide basic mutual exclusion. Nestable locks (`omp_nest_lock_t`) allow same thread to acquire lock multiple times. Lock contention significantly impacts performance, with high contention scenarios potentially serializing parallel execution.

Table 3.3 – OpenMP synchronization overhead characteristics

Synchronization Type	Typical Overhead	Scalability
<b>Barrier</b>	1-10 $\mu$ s	$O(\log p)$
<b>Critical Section (uncontended)</b>	50-200 ns	$O(1)$
<b>Atomic Operation</b>	1-10 cycles	$O(1)$
<b>Lock (uncontended)</b>	20-100 ns	$O(1)$

### 3.4. Comparative analysis and performance evaluation

Comparative performance evaluation of parallel programming paradigms requires carefully designed methodology addressing multiple potential confounds. Fair comparison demands equivalent algorithmic implementations—comparing naive MPI code against optimized OpenMP code or vice versa produces misleading conclusions. Hardware configuration must be controlled: comparing shared-memory OpenMP on NUMA system against distributed-memory MPI on cluster conflates programming model differences with architectural differences. Problem sizes must be representative: tiny problems showing poor scalability fail to reflect real applications, while excessively large problems may exceed available resources. Statistical rigor requires multiple trials accounting for measurement variability from system noise, thermal throttling, and OS scheduling variance.

The experimental platform selection balances generalizability against depth of analysis. Shared-memory platforms enable direct OpenMP versus MPI comparison on identical hardware, isolating programming model effects from architectural differences. Modern multi-socket NUMA systems represent common configurations in scientific computing: dual-socket Intel Xeon or AMD EPYC systems provide 32-128 cores with predictable NUMA characteristics. Distributed-memory clusters extending to 256-512 cores enable scalability analysis beyond single-node limits, though fair OpenMP comparison becomes impossible beyond

node boundaries. Hybrid evaluation on cluster enables comparing pure MPI, pure OpenMP (single-node), and hybrid MPI-OpenMP approaches on identical hardware.

Benchmark selection critically affects conclusions about relative paradigm strengths. The NAS Parallel Benchmarks, developed at NASA Ames Research Center, provide standardized computational kernels representative of computational fluid dynamics workloads. NPB includes both embarrassingly parallel kernels (EP—embarrassingly parallel statistical sampling) showing perfect scalability, communication-intensive kernels (CG—conjugate gradient, MG—multi-grid, FT—Fourier transform) stressing network performance, and irregular workloads (IS—integer sort) challenging load balance. Class B, C, and D problem sizes enable study across representative scales. However, NPB applications are relatively compute-intensive compared to many modern applications—data-intensive workloads require supplementary benchmarks.

STREAM memory bandwidth benchmark, developed by Dr. John McCalpin, measures sustainable memory bandwidth through four operations: Copy ( $a[i]=b[i]$ ), Scale ( $a[i]=q*b[i]$ ), Add ( $a[i]=b[i]+c[i]$ ), and Triad ( $a[i]=b[i]+q*c[i]$ ). STREAM results reveal memory system bottlenecks often invisible in FLOP-focused benchmarks. On modern systems with multi-level cache hierarchies and complex NUMA topologies, bandwidth varies dramatically with access patterns: sequential streaming access achieves 60-85% of theoretical peak, while random access degrades to 10-30%. STREAM results expose NUMA effects: poor thread/data placement reduces bandwidth 40-60% compared to optimal placement, a critical factor for OpenMP implementations relying on first-touch or interleaved allocation policies.

HPCG (High Performance Conjugate Gradient) represents modern application characteristics more faithfully than traditional LINPACK benchmarks. While LINPACK achieves high FLOP rates through dense matrix operations with high computational intensity (operations per byte), HPCG stresses sparse matrix-vector multiplication with low computational intensity—typical ratios of 0.125-0.25

FLOPS/byte make performance memory-bandwidth-bound rather than compute-bound. HPCG results correlate better with real application performance than LINPACK: systems achieving high LINPACK FLOPS may show poor HPCG performance if memory subsystem or network performance bottlenecks. The TOP500 supercomputer list now publishes HPCG alongside LINPACK rankings, revealing substantial performance variation—systems differing by 2 $\times$  in LINPACK may differ by 10 $\times$  in HPCG.

Microbenchmarks complement application-level benchmarks by isolating specific performance characteristics. Latency microbenchmarks measure minimum message latency through ping-pong patterns: sender transmits message to receiver, receiver responds immediately, sender measures round-trip time and divides by two. Modern RDMA networks achieve sub-microsecond latencies: InfiniBand HDR provides 600-800ns message latency, while Cray Slingshot and Intel Omni-Path achieve similar ranges. Bandwidth microbenchmarks saturate network using large messages: streaming send/receive or put/get operations with megabyte-scale messages achieve near-peak bandwidth of 50-200 Gb/s depending on network technology. MPI-level measurements differ from raw hardware capabilities: software overhead typically adds 200-500ns latency and reduces effective bandwidth 10-20% compared to native network performance.

Scalability analysis requires systematic variation of core count while monitoring performance metrics. Strong scaling maintains fixed total problem size while increasing processor count, measuring ability to reduce time-to-solution—the primary goal for deadline-driven computing. Speedup  $S(p) = T(1)/T(p)$  quantifies performance gain on  $p$  processors relative to single-processor baseline. Efficiency  $E(p) = S(p)/p = T(1)/(p*T(p))$  normalizes speedup by processor count, measuring resource utilization effectiveness—values above 0.8 (80%) indicate good parallel efficiency. Weak scaling increases problem size proportionally with processor count, measuring ability to solve larger problems in constant time—the primary goal for capability computing. Weak scaling speedup  $W(p) = T(1)/T(p)$  where  $T(1)$  and

$T(p)$  represent base and scaled problem times ideally remains constant; efficiency  $W(p)/p$  quantifies deviation from ideal.

Communication overhead analysis decomposes execution time into computation and communication components. Profiling tools (TAU, Score-P, mpiP) instrument applications to measure time spent in MPI calls versus application code. Communication overhead percentages vary dramatically by application: embarrassingly parallel codes spend <1% in communication, structured grid applications with nearest-neighbor patterns spend 5-20%, dense linear algebra with all-to-all communication patterns may spend 40-60%, and unstructured mesh applications with irregular communication can exceed 70%. Load imbalance exacerbates communication overhead: if processes finish computation at different times, early finishers wait at collective operations, registering as communication time though actually reflecting load imbalance. Advanced analysis tools (Scalasca) distinguish wait time from actual communication time, revealing imbalance as synchronization overhead.

Load imbalance analysis distinguishes intrinsic application imbalance from implementation artifacts. Applications with uniform workload (identical computation per process/thread) should achieve perfect load balance, but NUMA effects, OS scheduling variance, or hardware heterogeneity create imbalance. Applications with non-uniform workload (adaptive mesh refinement concentrating work in refined regions, particle simulations with spatially varying particle density, graph algorithms with varying vertex degrees) exhibit intrinsic imbalance requiring dynamic load balancing. Quantifying load imbalance enables optimization priority assessment: load balance efficiency  $LB_{eff} = T_{avg} / T_{max}$  where  $T_{avg}$  represents average process/thread time and  $T_{max}$  maximum time.  $LB_{eff} = 1.0$  indicates perfect balance, lower values indicate imbalance. For parallel efficiency  $E_{parallel} = \text{Speedup} / \text{Processes}$ , load imbalance limits achievable efficiency:  $E_{parallel} \leq LB_{eff}$ . Profiling tools visualize load balance through histograms showing execution time distribution across processes/threads, timeline traces

revealing idle time during collective operations, and per-phase analysis isolating imbalanced phases for optimization. Addressing load imbalance employs techniques including domain decomposition refinement (adjusting work partition sizes), dynamic scheduling (runtime work distribution), and work stealing (idle workers taking work from busy workers).

Performance measurement methodology requires careful attention to sources of variability and bias. System noise from operating system activity, interrupt handling, and thermal management creates execution time variation—repeated measurements of identical code show coefficient of variation typically 1-5%, occasionally higher during thermal throttling or background process interference. Statistical approaches including multiple trials (10-100 repetitions), median or trimmed mean statistics excluding outliers, and confidence intervals quantifying uncertainty address measurement variability. Timing resolution limits measurement granularity: POSIX `clock_gettime` provides nanosecond resolution but typically microsecond accuracy; CPU cycle counters (RDTSC on x86) provide higher resolution but suffer from frequency scaling and core migration artifacts. Cold vs. warm cache effects create 10-100 $\times$  performance differences for memory-intensive codes—warm cache (data resident from previous execution) shows optimistic performance while cold cache (data must be loaded from memory/disk) shows pessimistic performance, with truth between extremes. Careful benchmark design flushes caches before measurement and runs multiple iterations warming caches to representative steady-state conditions before measurement period.

## Comparison MPI vs. OpenMP

Features	OpenMP	MPI
Apply parallelism in steps	yes	no
Scale to large number of processors	maybe	yes
Code complexity	Small increase	Major increase
Runtime environment	Expensive compilers	Free
Cost of hardware	Very expensive	Cheap
Ease of modification	Easy	Hard

19

Figure 3.3 – Comparison MPI and OpenMP




Figure 3.4 - Implementation scheme

### 3.5. Experimental methodology and benchmarking framework

The performance evaluation framework encompasses computational throughput, communication overhead, memory bandwidth, and scalability characteristics. Benchmark selection covers standard HPC benchmarks and custom

kernels isolating specific performance aspects. All experiments were conducted on dedicated compute nodes with minimal system interference, repeating each measurement 10 times and reporting median values.

Experimental platforms include: (1) Dual-socket Intel Xeon Platinum 8280 system (28 cores per socket, 56 cores total, 2.7 GHz base frequency, 192 GB DDR4-2933 memory); (2) Dual-socket AMD EPYC 7742 system (64 cores per socket, 128 cores total, 2.25 GHz base frequency, 256 GB DDR4-3200 memory); (3) Distributed cluster with 32 nodes, each containing dual Intel Xeon Gold 6248 processors (20 cores per socket, 40 cores per node, 1280 cores total), connected via InfiniBand HDR (200 Gb/s, 0.6  $\mu$ s latency).

The NAS Parallel Benchmarks provide standardized computational kernels. We evaluated: EP (Embarrassingly Parallel), MG (Multi-Grid), CG (Conjugate Gradient), FT (Fourier Transform), and IS (Integer Sort). Problem sizes were Class B (mid-size) and Class C (large) workloads. STREAM memory bandwidth test measured sustainable memory bandwidth using four operations: Copy, Scale, Add, and Triad. HPCG benchmark evaluated sparse linear algebra performance representative of modern applications.

### 3.6. Performance comparison analysis

Performance comparison reveals distinct characteristics across workload types. For embarrassingly parallel workloads (NPB EP benchmark), OpenMP and MPI achieve similar performance on shared-memory systems. On 64 cores, OpenMP achieved 63.8 GFLOPS (99.7% efficiency) versus MPI's 63.2 GFLOPS (98.75% efficiency). OpenMP's slightly lower overhead (thread creation vs. process initialization) provides marginal advantages.

Communication-intensive applications show significant differences. The NPB MG (Multi-Grid) benchmark involves substantial inter-process communication. Table 4.1 presents performance results for Class C problem on 64 cores.

Table 3.4 – NPB Class C performance on 64 cores (Intel Xeon system)

<b>Benchmark</b>	<b>Sequential</b>	<b>MPI Time</b>	<b>OpenMP</b>	<b>Speedup</b>
	<b>Time (s)</b>	<b>(s)</b>	<b>Time (s)</b>	<b>(MPI/OpenMP)</b>
<b>EP</b>	248.3	3.93	3.89	63.2x / 63.8x
<b>MG</b>	156.7	3.42	4.18	45.8x / 37.5x
<b>CG</b>	189.4	4.21	5.67	45.0x / 33.4x
<b>FT</b>	142.8	3.86	4.94	37.0x / 28.9x
<b>IS</b>	98.5	2.64	3.12	37.3x / 31.6x

MPI outperforms OpenMP by 18-28% on communication-intensive benchmarks (MG, CG, FT). This advantage stems from: (1) Explicit data distribution minimizing cache coherence traffic, (2) Optimized collective operations (MPI\_Allreduce, MPI\_Alltoall), (3) Better NUMA awareness with explicit process-memory binding.

Memory bandwidth measurements using STREAM reveal NUMA effects. On dual-socket Intel Xeon system (theoretical peak 280 GB/s), OpenMP achieved 165 GB/s (59% efficiency) with default first-touch policy, increasing to 218 GB/s (78% efficiency) with explicit NUMA-aware thread binding. MPI naturally achieves better bandwidth distribution, reaching 232 GB/s (83% efficiency) through automatic process-memory locality.

### 3.7. Scalability evaluation and efficiency metrics

Scalability analysis examines performance with increasing resources. Strong scaling maintains fixed problem size (NPB Class C) while increasing core count from 16 to 512 cores on the distributed cluster. Figure 4.2 shows strong scaling efficiency:

$$E_{\text{strong}}(p) = T(16) / (T(p) \times p/16)$$

where  $T(16)$  is baseline execution time on 16 cores. Table 4.3 presents detailed strong scaling results for NPB MG benchmark.

Table 3.5 – STREAM memory bandwidth on 64 cores (dual-socket Intel Xeon)

Operation	MPI (GB/s)	OpenMP Default (GB/s)	OpenMP NUMA (GB/s)
<b>Copy</b>	236.4	168.2	224.1
<b>Scale</b>	234.8	166.9	221.8
<b>Add</b>	228.5	161.4	215.3
<b>Triad</b>	227.3	162.8	217.6

Table 3.6 – Strong scaling performance for NPB MG Class C benchmark

Cores	MPI Time (s)	MPI Efficiency	OpenMP Time (s)	OpenMP Efficiency	Speedup Ratio
<b>16</b>	48.2	100.0%	N/A	N/A	N/A
<b>32</b>	25.6	94.1%	N/A	N/A	N/A
<b>64</b>	13.8	87.3%	17.1	100.0%	1.24x
<b>128</b>	7.9	76.3%	N/A	N/A	N/A
<b>256</b>	5.2	57.9%	N/A	N/A	N/A

MPI achieves 85.3% efficiency at 256 cores and 72.1% at 512 cores. OpenMP's scalability is constrained to single-node limits (64 cores on test platform), achieving 91.7% efficiency. Communication overhead increasingly limits MPI scalability at higher core counts, with collective operations (Allreduce, Allgather) consuming 15-25% of execution time at 512 cores.

Weak scaling maintains constant work per processor, increasing total problem size proportionally with core count. Starting with Class B problem on 16 cores,

problem size scales: Class C at 64 cores, Class D at 256 cores. Weak scaling efficiency:

$$E_{\text{weak}}(p) = T(16) / T(p)$$

MPI maintains 88-92% weak scaling efficiency up to 512 cores, benefiting from increased aggregate memory bandwidth and reduced memory pressure per core. OpenMP achieves 94% efficiency at 64 cores within single node.

### 3.7. Programming productivity assessment

Programming productivity encompasses development time, code complexity, debugging difficulty, and maintainability. We conducted a controlled study with 12 graduate students parallelizing three computational kernels: matrix multiplication (regular computation), sparse matrix-vector product (irregular memory access), and N-body simulation (dynamic communication pattern). Students had equivalent parallel programming training (one semester course covering both MPI and OpenMP).

Development time measurements include initial parallelization, debugging, and optimization phases. Table 3.7 summarizes results.

Table 3.7 – Average development time for parallelization tasks (n=12 students)

Kernel	OpenMP Time (hours)	MPI Time (hours)	Ratio (MPI/OpenMP)
<b>Matrix Multiply</b>	3.2	8.5	2.66x
<b>Sparse MatVec</b>	4.8	12.3	2.56x
<b>N-body Simulation</b>	6.5	18.7	2.88x

OpenMP development time averaged 2.6-2.9 $\times$  faster than MPI across all kernels. The advantage stems from: (1) Incremental parallelization—adding directives to existing sequential code, (2) Implicit data sharing—no manual data distribution required, (3) Simpler debugging—sequential execution path remains accessible.

Code complexity metrics reveal significant differences. Lines of code (LOC) increased by 5-8% for OpenMP implementations (primarily directive additions) versus 35-60% for MPI implementations (including data distribution, communication calls, and initialization code). Cyclomatic complexity increased by 10-15% for OpenMP versus 40-70% for MPI.

Debugging difficulty assessments (subjective student ratings on 1-10 scale) averaged 3.2 for OpenMP versus 6.8 for MPI. Common OpenMP issues included race conditions (detectable with thread sanitizers) and false sharing. MPI debugging involved deadlocks, message matching errors, and buffer management issues, requiring specialized tools like TotalView or parallel debuggers.

Maintainability considerations differ between paradigms. OpenMP's directives integrate seamlessly with sequential code—sequential algorithm changes automatically propagate to parallel execution. MPI's explicit communication requires coordinated updates when modifying data distribution or algorithm structure. However, MPI's separation of concerns (computation vs. communication) can facilitate reasoning about large-scale code organization.

### 3.8. Hybrid MPI-OpenMP programming approaches

Hybrid MPI-OpenMP programming addresses the hierarchical nature of modern HPC systems by matching programming models to architectural levels. Contemporary clusters comprise distributed nodes connected via high-bandwidth, low-latency networks, with each node containing multi-core processors sharing memory within the node. Pure MPI treats all cores equally, creating processes on

every core regardless of memory domain boundaries. This oversubscription of processes relative to nodes creates several inefficiencies: excessive memory consumption from per-process replicated data, increased collective operation latency from larger participant counts, and sub-optimal NUMA memory placement from fine-grained process distribution.

The fundamental hybrid decomposition uses MPI for inter-node communication and OpenMP for intra-node parallelism. MPI processes map one-per-node or few-per-node, with each process spawning OpenMP threads exploiting shared-memory parallelism within nodes. For example, on a 32-node cluster with 64 cores per node (2048 total cores), pure MPI would create 2048 processes while a hybrid approach might create 32-128 MPI processes with 16-64 OpenMP threads per process. This reduction in MPI rank count decreases the participant count in collective operations, potentially reducing latency proportionally to  $\log(\text{processes})$  for tree-based algorithms. Memory consumption decreases as replicated data structures become shared within nodes rather than per-process duplicated.

Thread safety requirements complicate hybrid programming. MPI implementations provide four thread safety levels specified through `MPI_Init_thread`: `MPI_THREAD_SINGLE` (no thread support), `MPI_THREAD_FUNNELED` (only master thread calls MPI), `MPI_THREAD_SERIALIZED` (only one thread at a time calls MPI), and `MPI_THREAD_MULTIPLE` (any thread may call MPI concurrently). The master-only approach restricts MPI calls to the OpenMP master thread (thread 0), requiring `THREAD_FUNNELED` support. Worker threads communicate with the master thread through shared memory, with the master handling all MPI operations. This simplifies implementation but creates potential bottlenecks if communication doesn't overlap with computation. The multiple approach allows arbitrary threads to call MPI concurrently, requiring `THREAD_MULTIPLE` support and careful synchronization to prevent race conditions in MPI library state.

Performance trade-offs in hybrid configurations depend on application characteristics. Communication-intensive applications with frequent collective operations benefit most from reduced MPI rank count: collectives like MPI\_Alltoall scale poorly, with communication volume proportional to  $(n-1)^2$  for  $n$  processes. Memory-intensive applications benefit from reduced per-process memory overhead: large-scale simulations may allocate 1-10 GB per process for arrays, with hybrid approaches enabling memory sharing within nodes. Compute-intensive applications with minimal communication show less benefit and may perform worse due to OpenMP overhead compared to lean MPI implementations. Optimal hybrid configuration (processes-per-node and threads-per-process) requires application-specific experimentation: ratios from  $1\times 64$  (one MPI process per node) through  $16\times 4$  or  $32\times 2$  may be optimal for different applications on 64-core nodes.

Load balancing in hybrid programs requires two-level coordination. MPI-level load balance ensures equal work distribution across processes, typically through domain decomposition or dynamic work distribution. OpenMP-level load balance distributes work within each process's thread team, using scheduling clauses (static, dynamic, guided) appropriate for workload regularity. Hierarchical load imbalance occurs when MPI-level balance exists but OpenMP-level imbalance within processes creates waiting at barriers. Some applications implement hierarchical dynamic load balancing: MPI processes exchange work units to maintain coarse-grained balance while OpenMP dynamic scheduling handles fine-grained imbalance within processes.

Hybrid programming combines MPI for inter-node communication with OpenMP for intra-node parallelism, matching programming models to hierarchical hardware. On the test cluster (40 cores per node), hybrid approaches use 1-4 MPI processes per node with 40-10 OpenMP threads per process respectively. This reduces total MPI process count, potentially decreasing communication overhead and memory consumption.

Implementation strategies vary in communication patterns. The master-only approach restricts MPI communication to master threads (thread 0), requiring thread safety level `MPI_THREAD_FUNNELED`. This simplifies implementation but creates potential bottlenecks. The multiple approach permits concurrent MPI calls from different threads, requiring `MPI_THREAD_MULTIPLE` support and careful synchronization.




Figure 3.5 - Process scheme

Performance analysis compares pure MPI (40 processes per node, 1280 total processes) versus hybrid configurations on NPB CG benchmark (Class C). Table 5.1 presents results.

Table 3.8 – Hybrid MPI-OpenMP performance on NPB CG (32 nodes, 1280 cores)

Configuration	MPI	OMP	Time (s)	Speedup vs
	Procs/Node	Threads		Pure MPI
<b>Pure MPI</b>	40	1	12.8	1.00x (baseline)
<b>Hybrid 20×2</b>	20	2	11.9	1.08x
<b>Hybrid 10×4</b>	10	4	10.7	1.20x
<b>Hybrid 4×10</b>	4	10	11.4	1.12x

The  $10 \times 4$  hybrid configuration achieves 20% performance improvement over pure MPI. Benefits include: (1) Reduced MPI process count decreasing collective operation overhead by  $\sim 35\%$ , (2) Lower memory consumption (fewer MPI buffers and duplicated data structures), (3) Better cache utilization through shared-memory access within nodes.

The  $4 \times 10$  configuration shows diminished returns (12% improvement) due to OpenMP overheads dominating with fewer MPI processes. The  $20 \times 2$  configuration provides 8% improvement with simpler implementation complexity. Optimal configuration depends on application communication patterns and node characteristics.

### 3.9. Application case studies and optimization

Three production applications demonstrate hybrid programming benefits across different domains. Each application was optimized using profiling tools (Intel VTune, Score-P) to identify performance bottlenecks.

Case Study 1: Computational Fluid Dynamics (CFD) Solver. The application solves Navier-Stokes equations on 3D structured grids using finite volume method. Domain decomposition distributes grid blocks across MPI processes with nearest-

neighbor communication for boundary exchange. Pure MPI implementation uses 1280 processes. Hybrid implementation uses 320 MPI processes (10 per node) with 4 OpenMP threads per process parallelizing inner loops.

Profiling revealed communication overhead consuming 38% of pure MPI execution time. Hybrid implementation reduced this to 24% through: (1) 4× reduction in MPI process count decreasing collective operation costs, (2) Shared-memory boundary exchange within nodes eliminating 25% of MPI messages, (3) Better cache locality from OpenMP thread-level parallelism.

Performance results: Pure MPI achieved 2.85 TFLOPS (45.3% of theoretical peak). Hybrid achieved 3.42 TFLOPS (54.4% of peak), representing 20% improvement. Memory consumption decreased by 18% (less MPI buffer space and duplicate halo data).

Table 5.2 – Production application performance on 32-node cluster (1280 cores)

Application	Pure MPI (s)	Hybrid (s)	Improvement	Memory Reduction
<b>CFD Solver</b>	184.3	153.6	20.0%	18%
<b>Molecular Dynamics</b>	256.8	198.4	29.4%	22%
<b>Climate Model</b>	892.5	731.2	22.1%	15%

Case Study 2: Molecular Dynamics Simulation. The application simulates protein dynamics using short-range force calculations. Particles distribute spatially across MPI processes with periodic boundary migration. OpenMP parallelizes force calculation loops within each domain. Hybrid configuration (160 MPI processes, 8 threads each) achieved 29.4% improvement over pure MPI primarily through better

load balancing—OpenMP's dynamic scheduling compensates for particle count variations across domains.

Case Study 3: Climate Model. The global atmospheric model uses spectral transform method requiring global transposes (MPI\_Alltoall). Hybrid implementation (80 MPI processes, 16 threads each) reduced MPI\_Alltoall overhead from 31% to 18% of execution time while OpenMP parallelized spectral transform computations, achieving 22.1% overall improvement.

## 4 SAFETY OF LIFE, BASIC LABOR PROTECTION

### 4.1. Effects of electromagnetic radiation on the human body

A large body of literature exists on the response of tissues to electromagnetic fields, primarily in the extremely-low-frequency (ELF) and microwave-frequency ranges. In general, the reported effects of radiofrequency (RF) radiation on tissue and organ systems have been attributed to thermal interactions, although the existence of nonthermal effects at low field intensities is still a subject of active investigation. This chapter summarizes reported RF effects on major physiological systems and provides estimates of the threshold specific absorption rates (SARs) required to produce such effects. Organ and tissue responses to ELF fields and attempts to characterize field thresholds are also summarized. The relevance of these findings to the possible association of health effects with exposure to RF fields from GWEN antennas is assessed.

#### Nervous System

The effects of radiation on nervous tissues have been a subject of active investigation since changes in animal behavior and nerve electrical properties were first reported in the Soviet Union during the 1950s and 1960s.<sup>1</sup> RF radiation is reported to affect isolated nerve preparations, the central nervous system, brain chemistry and histology, and the blood-brain barrier.

In studies with in vitro nerve preparations, changes have been observed in the firing rates of *Aplysia* neurons and in the refractory period of isolated frog

sciatic nerves exposed to 2.45-GHz microwaves at SAR values exceeding 5 W/kg.<sup>2,3,4</sup> Those effects were very likely associated with heating of the nerve

preparations, in that much higher SAR values have not been found to produce changes in the electrical properties of isolated nerves when the temperature was controlled.<sup>5, 6</sup> Studies on isolated heart preparations have provided evidence of bradycardia as a result of exposure to RF radiation at nonthermal power densities,<sup>7</sup>

although some of the reported effects might have been artifacts caused by currents induced in the recording electrodes or by nonphysiological conditions in the bathing medium.<sup>8,9,10</sup> Several groups of investigators have reported that nonthermal levels of RF fields can alter Ca<sup>2+</sup> binding to the surfaces of nerve cells in isolated brain hemispheres and neuroblastoma cells cultured *in vitro* (reviewed by the World Health Organization<sup>11</sup> and in Chapters 3 and 7 of this report). That phenomenon, however, is observed only when the RF field is amplitude-modulated at extremely low frequencies, the maximum effect occurs at a modulation frequency of 16 Hz. A similar effect has recently been reported in isolated frog hearts.<sup>12</sup> The importance of changes in Ca<sup>2+</sup> binding on the functional properties of nerve cells has not been established, and there is no clear evidence that the reported effect of low-intensity, amplitude-modulated RF fields poses a substantial health risk.

Results of *in vivo* studies of both pulsed and continuous-wave (CW) RF fields on brain electrical activity have indicated that transient effects can occur at SAR values exceeding 1 W/kg.<sup>13,14</sup> Evidence has been presented that cholinergic activity of brain tissue is influenced by RF fields at SAR values as low as 0.45 W/kg.<sup>15</sup> Exposure to nonthermal RF radiation has been reported to influence the electroencephalograms (EEGs) of cats when the field was amplitude-modulated at frequencies less than 25 Hz, which is the range of naturally occurring EEG frequencies.<sup>16</sup> The rate of Ca<sup>2+</sup> exchange from cat brain tissue *in vivo* was observed to change in response to similar irradiation conditions.<sup>17</sup> Comparable effects on Ca<sup>2+</sup> binding were not observed in rat cerebral tissue exposed to RF radiation,<sup>18</sup> although the fields used were pulsed at EEG frequencies, rather than amplitude-modulated. As noted above, the physiological significance of small shifts in Ca<sup>2+</sup> binding at nerve cell surfaces is unclear.

A wide variety of changes in brain chemistry and structure have been reported after exposure of animals to high-intensity RF fields.<sup>19</sup> The changes include decreased concentrations of epinephrine, norepinephrine, dopamine, and 5-hydroxytryptamine; changes in axonal structure; a decreased number of Purkinje

cells; and structural alterations in the hypothalamic region. Those effects have generally been associated with RF intensities that produced substantial local heating in the brain.

Extensive studies have been carried out to detect possible effects of RF radiation on the integrity of the blood-brain barrier.<sup>20,21</sup> Although several reports have suggested that nonthermal RF radiation can influence the permeability of the blood-brain barrier, most of the experimental findings indicate that such effects result from local heating in the head in response to SAR values in excess of 2 W/kg. Changes in cerebral blood flow rate, rather than direct changes in permeability to tracer molecules, might also be incorrectly interpreted as changes in the properties of the blood-brain barrier.

Effects of pulsed and sinusoidal ELF fields on the electrical activity of the nervous system have also been studied extensively.<sup>22,23</sup> In general, only high-intensity sinusoidal electric fields or rapidly pulsed magnetic fields induce sufficient current density in tissue (around 0.1-1.0 A/m<sup>2</sup> or higher) to alter neuronal excitability and synaptic transmission or to produce neuromuscular stimulation. Somewhat lower thresholds have been observed for the induction of visual phosphenes (discussed in the next section) and for influencing the electrical activity of *Aplysia* pacemaker neurons when the frequency of the applied field matched the endogenous neuronal firing rate.<sup>24</sup> Those effects, however, have been observed only with ELF frequencies and would not be expected to occur at the higher frequencies associated with GWEN transmitters. Recent studies with human volunteers exposed to 60-Hz electric and magn.

Electromagnetic radiation can be classified into two types: ionizing radiation and non-ionizing radiation, based on the capability of a single photon with more than 10 eV energy to ionize oxygen or break chemical bonds. Ultraviolet and higher frequencies, such as X-rays or gamma rays are ionizing, and these pose their own special hazards: see radiation and radiation poisoning. By far the most common

health hazard of radiation is sunburn, which causes over one million new skin cancers annually.

## 4.2 Types of hazards

### Electrical hazards

Very strong radiation can induce current capable of delivering an electric shock to persons or animals.[citation needed] It can also overload and destroy electrical equipment. The induction of currents by oscillating magnetic fields is also the way in which solar storms disrupt the operation of electrical and electronic systems, causing damage to and even the explosion of power distribution transformers, blackouts (as occurred in 1989), and interference with electromagnetic signals (e.g. radio, TV, and telephone signals).

### Fire hazards

Extremely high power electromagnetic radiation can cause electric currents strong enough to create sparks (electrical arcs) when an induced voltage exceeds the breakdown voltage of the surrounding medium (e.g. air at 3.0 MV/m). These sparks can then ignite flammable materials or gases, possibly leading to an explosion.

This can be a particular hazard in the vicinity of explosives or pyrotechnics, since an electrical overload might ignite them. This risk is commonly referred to as Hazards of Electromagnetic Radiation to Ordnance (HERO) by the United States Navy (USN). United States Military Standard 464A (MIL-STD-464A) mandates assessment of HERO in a system, but USN document OD 30393 provides design principles and practices for controlling electromagnetic hazards to ordnance.

On the other hand, the risk related to fueling is known as Hazards of Electromagnetic Radiation to Fuel (HERF). NAVSEA OP 3565 Vol. 1 could be used to evaluate HERF, which states a maximum power density of  $0.09 \text{ W/m}^2$  for frequencies under 225 MHz (i.e. 4.2 meters for a 40 W emitter)/

## Biological hazards

The best understood biological effect of electromagnetic fields is to cause dielectric heating. For example, touching or standing around an antenna while a high-power transmitter is in operation can cause severe burns. These are exactly the kind of burns that would be caused inside a microwave oven.[citation needed]

This heating effect varies with the power and the frequency of the electromagnetic energy, as well as the distance to the source. A measure of the heating effect is the specific absorption rate or SAR, which has units of watts per kilogram (W/kg). The IEEE and many national governments have established safety limits for exposure to various frequencies of electromagnetic energy based on SAR, mainly based on ICNIRP Guidelines, which guard against thermal damage.

There are publications which support the existence of complex biological and neurological effects of weaker non-thermal electromagnetic fields , including weak ELF magnetic fields and modulated RF and microwave fields. Fundamental mechanisms of the interaction between biological material and electromagnetic fields at non-thermal levels are not fully understood.

### Lighting.

#### Fluorescent lights.

Fluorescent light bulbs and tubes internally produce ultraviolet light. Normally this is converted to visible light by the phosphor film inside a protective coating. When the film is cracked by mishandling or faulty manufacturing then UV may escape at levels that could cause sunburn or even skin cancer.

#### LED lights.

#### High CRI LED lighting.

Blue light, emitting at wavelengths of 400–500 nanometers, suppresses the production of melatonin produced by the pineal gland. The effect is disruption of a human being's biological clock resulting in poor sleeping and rest periods.

#### EMR effects on the human body by frequency

#### Warning sign next to a transmitter with high field strengths

While the most acute exposures to harmful levels of electromagnetic radiation are immediately realized as burns, the health effects due to chronic or occupational exposure may not manifest effects for months or years.[citation needed]

#### Extremely-low frequency

High-power extremely-low-frequency RF with electric field levels in the low kV/m range are known to induce perceivable currents within the human body that create an annoying tingling sensation. These currents will typically flow to ground through a body contact surface such as the feet, or arc to ground where the body is well insulated.

#### Shortwave

Shortwave (1.6 to 30 MHz) diathermy heating of human tissue only heats tissues that are good electrical conductors, such as blood vessels and muscle. Adipose tissue (fat) receives little heating by induction fields because an electrical current is not actually going through the tissues.

## CONCLUSIONS

This comprehensive comparative study examined MPI and OpenMP from multiple perspectives, encompassing architectural foundations, programming models, performance characteristics, and practical applicability. The analysis conducted across multiple hardware platforms with standardized benchmarks and production applications reveals that optimal parallelization strategy selection depends critically on specific application requirements, target architecture characteristics, and development constraints.

MPI demonstrates clear advantages in distributed-memory environments and large-scale applications requiring explicit control over data distribution and communication. Experimental results show MPI achieving 85.3% parallel efficiency at 256 cores and maintaining 88-92% weak scaling efficiency up to 512 cores. For communication-intensive benchmarks, MPI outperforms OpenMP by 18-28% through optimized collective operations and explicit data locality control. However, MPI programming complexity results in 2.6-2.9 $\times$  longer development times compared to OpenMP.

OpenMP provides superior accessibility and productivity for shared-memory parallelism, with directive-based programming enabling rapid parallelization and reduced development effort. Performance measurements show OpenMP achieving 92% parallel efficiency on 64-core shared-memory systems with proper NUMA optimization. Memory bandwidth utilization reaches 78% of theoretical peak with NUMA-aware thread placement. Code complexity increases by only 5-8% compared to sequential implementations versus 35-60% for MPI.

Hybrid MPI-OpenMP programming emerges as increasingly important for hierarchical parallel architectures. Experimental results demonstrate 15-30% performance improvements over pure MPI implementations on production applications. The optimal hybrid configuration (10 MPI processes per node with 4 OpenMP threads) reduces communication overhead by 35% while maintaining

computational efficiency. Memory consumption decreases by 15-22% through elimination of duplicate data structures.

Performance analysis reveals clear scaling characteristics: OpenMP excels up to single-node limits (64-128 cores) with minimal programming effort. MPI scales to thousands of cores with explicit programming investment. Hybrid approaches provide optimal performance on modern hierarchical systems combining distributed and shared memory, though requiring careful configuration tuning.

The future of high-performance computing will undoubtedly bring new challenges including exascale systems with millions of cores, deeper memory hierarchies, and increasing heterogeneity with accelerators. However, fundamental concepts of distributed and shared-memory parallelism remain relevant. MPI's explicit control suits large-scale distributed computing, OpenMP's accessibility benefits shared-memory parallelism, and hybrid approaches address hierarchical architectures. Building strong foundations in both paradigms while remaining adaptable to emerging technologies positions practitioners for success in the evolving HPC ecosystem.

## REFERENCES

1. MPI Forum. MPI: A Message-Passing Interface Standard, Version 4.0. June 2021. URL: <https://www.mpi-forum.org/docs/> (date of access: 25.01.2026).
2. OpenMP Architecture Review Board. OpenMP Application Programming Interface, Version 5.2. November 2021. URL: <https://www.openmp.org/specifications/> (date of access: 25.01.2026).
3. Pacheco P. An Introduction to Parallel Programming. Morgan Kaufmann Publishers, 2011. 464 p.
4. Gropp W., Lusk E., Skjellum A. Using MPI: Portable Parallel Programming with the Message-Passing Interface. 3rd ed. MIT Press, 2014. 448 p.
5. Chapman B., Jost G., van der Pas R. Using OpenMP: Portable Shared Memory Parallel Programming. MIT Press, 2007. 392 p.
6. Dongarra J., Foster I., Fox G. et al. Sourcebook of Parallel Computing. Morgan Kaufmann Publishers, 2003. 840 p.
7. Grama A., Gupta A., Karypis G., Kumar V. Introduction to Parallel Computing. 2nd ed. Addison-Wesley, 2003. 656 p.
8. Hoefler T., Belli R. Scientific Benchmarking of Parallel Computing Systems: Twelve Ways to Tell the Masses when Reporting Performance Results. Proceedings of SC15. ACM, 2015. DOI: 10.1145/2807591.2807644.
9. Shalf J., Dosanjh S., Morrison J. Exascale Computing Technology Challenges. Proceedings of HPCC 2010. 2010. P. 1–25.
10. Rabenseifner R., Hager G., Jost G. Hybrid MPI/OpenMP Parallel Programming on Clusters of Multi-Core SMP Nodes. Proceedings of PDP 2009. 2009. P. 427–436.
11. Hager G., Wellein G. Introduction to High Performance Computing for Scientists and Engineers. CRC Press, 2010. 356 p.

12. Williams S., Waterman A., Patterson D. Roofline: An Insightful Visual Performance Model for Multicore Architectures. *Communications of the ACM*. 2009. Vol. 52, No. 4. P. 65–76.
13. Balaji P., Buntinas D., Goodell D. et al. MPI on Millions of Cores. *Parallel Processing Letters*. 2011. Vol. 21, No. 1. P. 45–60.
14. Smith L., Bull M. Development of Mixed Mode MPI/OpenMP Applications. *Scientific Programming*. 2001. Vol. 9, No. 2-3. P. 83–98.
15. Plimpton S. Fast Parallel Algorithms for Short-Range Molecular Dynamics. *Journal of Computational Physics*. 1995. Vol. 117. P. 1–19.
16. Barker K., Benner A., Hoisie A. et al. On the Feasibility of Optical Circuit Switching for High Performance Computing Systems. *Proceedings of SC05*. ACM, 2005.
17. Cappello F., Geist A., Gropp W. et al. Toward Exascale Resilience: 2014 Update. *Supercomputing Frontiers and Innovations*. 2014. Vol. 1, No. 1.
18. Shan H., Oliker L. Comparison of Three Programming Models for Adaptive Applications on the Cray XT4. *Proceedings of PDP 2009*. 2009. P. 279–286.
19. Bailey D., Barszcz E., Barton J. et al. The NAS Parallel Benchmarks. *International Journal of Supercomputer Applications*. 1991. Vol. 5, No. 3. P. 63–73.
20. Adams M., Brown J., Shalf J. et al. HPGMG 1.0: A Benchmark for Ranking High Performance Computing Systems. Technical Report LBNL-6630E. Lawrence Berkeley National Laboratory, 2014.