

IPDPS 2027 Call for Papers



**41st IEEE International Parallel &
Distributed Processing Symposium
June 1-5, 2027**

PLEASE NOTE:

Authors must register their paper and submit an abstract by **Thursday, October 1, 2026**.

Authors must then submit full versions of registered papers by **Thursday, October 8, 2026** (firm deadline).

All deadlines are end of day **ANYWHERE ON EARTH**.

Authors of accepted papers are expected to **present them in person** at the conference.

Before submitting, review the information under **WHAT/WHERE TO SUBMIT** below.

Authors are invited to submit manuscripts that present novel and impactful research in high performance computing (HPC) in parallel and distributed processing. R&D efforts focusing on emerging technologies, interdisciplinary work spanning multiple IPDPS focus areas, and novel open-source artifacts are welcome. Topics of interest include, but are not limited to, the following topic areas:

Algorithms:

This track focuses on algorithms for computational and data science in parallel and distributed computing environments (including cloud, edge, fog, distributed memory, and accelerator-based computing). Examples include structured and unstructured mesh and meshless methods, dense and sparse linear algebra computations, spectral methods, n-body computations, clustering, data mining, compression, and combinatorial algorithms such as graph and string algorithms. Also included in this track are algorithms that apply to tightly or loosely coupled systems, such as those supporting communication, synchronization, power management, distributed resource management, distributed data and transactions, and mobility. This track also includes novel algorithm designs and

implementations tailored to emerging architectures (such as ML/AI accelerators or quantum computing systems).

Applications:

This track focuses on real-world applications (combinatorial, scientific, engineering, data analysis, and visualization) that use parallel and distributed computing concepts (including cloud, edge, fog, distributed memory, and accelerator-based computing). Papers submitted to this track are expected to incorporate innovations that originate in specific target application areas and contribute novel methods and approaches that address core challenges in their scalable implementation. Contributions include the design, implementation, and evaluation of parallel and distributed applications, including implementations targeting emerging architectures (such as ML/AI accelerators or quantum computing systems) and application domain advances enabled by ML/AI.

Architecture:

This track focuses on existing and emerging architectures for high performance computing, including architectures for instruction-level and thread-level parallelism; manycore, multicore, accelerator, domain-specific, and special-purpose architectures (including ML/AI accelerators); reconfigurable architectures; memory technologies and hierarchies; volatile and non-volatile emerging memory technologies; co-design paradigms for processing-in-memory architectures; solid-state devices; exascale system designs; data center and warehouse-scale architectures; novel big data architectures; network and interconnect architectures; emerging technologies for interconnects; parallel I/O and storage systems; power-efficient and green computing systems; resilience, security, and dependable architectures; and emerging architectural principles for machine learning, approximate computing, quantum computing, neuromorphic, analog, and bio-inspired computing.

Machine Learning and Artificial Intelligence (ML/AI):

This track focuses on all areas of ML/AI **that are relevant to parallel and distributed computing**, including ML/AI training on resource-limited platforms (e.g., edge platforms); computational optimization methods for AI such as pruning, quantization, and knowledge distillation; parallel and distributed learning algorithms; energy-efficient methods for ML/AI; federated learning; agentic AI; design and implementation of ML/AI algorithms on parallel architectures (including distributed memory, GPUs, tensor cores, and emerging ML/AI accelerators); new ML/AI methods benefiting HPC applications or HPC system management; and design and development of ML/AI software pipelines (e.g., frameworks for distributed training, integration of compression into ML/AI pipelines, compiler techniques, and DSLs). Papers submitted to the ML/AI track should emphasize ML/AI innovations that are best reviewed by ML/AI experts. *Papers that emphasize core parallel computing topics applied to ML/AI workloads or applications that benefit from existing ML/AI tools should be submitted to the topic domain tracks rather than this ML/AI track.*

Measurements, Modeling, and Experiments:

This track focuses on experiments and performance-oriented studies in the practice of parallel and distributed computing. “Performance” may be construed broadly to include metrics related to time, energy, power, accuracy, and resilience, for instance. Topics include methods, experiments, and tools for measuring, evaluating, and/or analyzing performance for large-scale applications and systems; design and experimental evaluation of applications of parallel and distributed computing in simulation and analysis; experiments on the use of novel commercial or research accelerators and architectures, including quantum, neuromorphic, and other non-Von Neumann systems; innovations made in support of large-scale infrastructures and facilities; and experiences and methods for allocating and managing system and facility resources.

Programming Models, Compilers, and Runtime Systems:

This track covers topics ranging from the design of parallel programming models and paradigms to languages and compilers that support these models and paradigms, as well as runtime and middleware solutions. Software that is close to the application (rather than the bare hardware) but not specific to an application is included. Examples include frameworks targeting cloud and distributed systems; application frameworks for fault tolerance and resilience; software supporting data management, scalable data analytics, and similar workloads; and runtime systems for future novel computing platforms, including quantum, neuromorphic, and bio-inspired computing. This track also includes novel compiler techniques and frameworks that leverage machine learning methods.

System Software:

This track focuses on software that is close to hardware in parallel and distributed computing environments, including cloud, edge, fog, distributed memory, and accelerator-based computing. Topics include storage and I/O systems; system software for resource management, job scheduling, and energy efficiency; system software support for accelerators and heterogeneous HPC computing systems; interactions between the operating system, hardware, and other software layers; system software for data management; system software solutions for ML/AI workloads (e.g., energy-efficient software methods for ML/AI); system software support for fault tolerance and resilience; containers and virtual machines; specialized operating systems and related support for high-performance computing; system software for future novel computing platforms, including quantum, neuromorphic, and bio-inspired computing; and system software advances enabled by ML/AI.

WHAT/WHERE TO SUBMIT

Abstracts of at most 500 words **must** be submitted **by October 1, 2026 AOE**. Manuscripts must be submitted **by October 8, 2026 AOE**; to ensure fairness, no extensions will be given.

- Submitted manuscripts may not exceed ten (10) single-spaced double-column pages using 10-point size font on 8.5x11-inch pages (IEEE conference style), including all figures and tables.
- There is no page limit for references. Each reference cited must be complete, including all author names and a DOI (if available) or a direct webpage link.
- No supplementary sections or appendices are allowed beyond the stated page limit during the paper submission phase. Accepted papers (after the notifications) are required to submit a reproducibility appendix. More details are in the Reproducibility Initiative section below.
- The program committee will use a double-anonymous review process. Submitted manuscripts should not include author names and affiliations or otherwise disclose the identity of the authors due to the double-anonymous review process.
- The number of submissions allowed from the same author is limited to eight (8).
- **Submissions not complying with these requirements will be subject to administrative rejection.**

The IEEE conference-style templates for MS Word and LaTeX provided by IEEE eXpress Conference Publishing are available for download. [See the latest versions here](#). Files must be submitted by following the instructions at the [IPDPS 2027 Submission Site \(powered by Linklings—https://ssl.linklings.net/conferences/ipdps\)](https://ssl.linklings.net/conferences/ipdps). Authors must select a “primary” track for the submission; the primary track is the area most related to the paper’s contributions. An optional “secondary” track may also be selected.

REVIEW OF MANUSCRIPTS

All submitted manuscripts will be reviewed by the Program Committee under a double-anonymous, two-round review process. Submissions will be judged on originality, correctness, technical strength, significance, demonstrated or potential impact, quality of presentation, and interest and relevance to the conference.

Submitted manuscripts must NOT have appeared in or be under consideration for another conference, workshop, or journal. Submissions violating this requirement will face administrative rejection at any phase of the review process.

A high-quality submission should articulate its contributions in multiple aspects:

1. *Motivation*. Clearly state the paper’s objective and provide strong support to motivate the specific problem the submission addresses.

2. *Limitations of state-of-art approaches.* Unambiguously discuss and distinguish the paper's contributions from the most relevant and most recent prior works.
3. *Key insights and contributions.* Clearly articulate the major insights that enable the described approach and make it effective. Clearly specify the novelty of these insights, research challenges, and how the solutions advance the state of the art. Provide a list of key contributions, including flagship theoretical or experimental results and improvements over the existing solutions.
4. *Methodology.* Clearly specify key theoretical or experimental methodological details. Support the chosen methodological choices (e.g., cite and compare with the prior works that have evaluated their ideas using similar methodology). If a new methodology is adopted or theoretical assumptions differ from prior art, provide a detailed justification.
5. *Limitations of the proposed approach.* Articulate all significant limitations of the proposed approach and identify conclusions that are sensitive to assumptions made in the paper.

The Program Committee will assess submissions in the above aspects. Therefore, the authors should make these aspects clear when articulating their contributions.

Authors will have the opportunity to respond to the reviewers' questions and provide clarifications before the program committee makes their first-round decisions. Some submissions may not be invited to submit a response/rebuttal; these submissions will be notified with an "early-reject" decision by **November 30, 2026**. The response/rebuttal from the authors is due on **December 3, 2026**.

First-round decisions – "accept," "revise," or "reject" – will be sent by **December 18, 2026**. Authors of papers in the "revise" category will have the opportunity to submit a new version of their papers that address reviewers' comments. The revised submission and a cover letter explaining changes are due on **January 18, 2027**. An ensuing review will then provide decisions of "accept" or "reject"; papers will be rejected if the reviewers assess that the issues they raised were not satisfactorily addressed.

- Notification of final decisions will be sent by **February 1, 2027**
- Camera-ready papers are due on **February 19, 2027**

Questions may be sent to pc2027@ipdps.org.

Best Paper Award

The program committee will select a small set of papers as Best Paper finalists. One paper will be named the Best Paper.

Best Open-Source Contribution Award (for “Accepted papers”)

IPDPS welcomes submissions with technical contributions of open-source tools and dataset artifacts relevant to the parallel and distributed computing community. The authors of **accepted papers** may identify their submissions for consideration for the Best Open-Source Contribution award. A link will be provided on Linklings to submit the details of the artifacts. A dedicated open-source tool and dataset artifacts committee will evaluate such papers. A small set of such papers will be identified as Best Open-Source Contribution finalists. One paper will be recognized with the Best Open-Source Contribution Award.

The two award categories are not exclusive; a paper can be nominated for both the Best Paper award and Best Open-Source Contribution award.

REPRODUCIBILITY INITIATIVE (for “Accepted papers”)

To produce a standardized, long-lasting impact, IPDPS is requiring a computational result reproducibility appendix. This appendix should describe the processes used to obtain the computational results and should be appended to the accepted papers. The authors can clearly mark which results are likely to be reproducible and which may not be. The appendix will be due after acceptance. **More information can be found [here](#).**

ArXiv Submission Policy

Having an arXiv paper does not prohibit authors from submitting a paper to IPDPS 2027. arXiv papers are not peer-reviewed and not considered as formal publications; hence, they do not count as prior work. Authors are not expected to compare against arXiv papers that have not formally appeared in conference or journal proceedings. Authors must follow the double-anonymous submission guidelines even if a submitted paper is already on arXiv. Authors are encouraged to use preventive measures to reduce the chances of accidental breach of anonymity (e.g., use a different title in the submission or not upload/revise the arXiv version during the review period). Authors should not direct reviewers to arXiv versions of the paper; in their evaluations, the reviewers will consider only the material in the submitted paper.

Guidance on Artificial Intelligence (AI)-Generated Text

Tools like Grammarly or other AI assistants may be used to improve the submission presentation. However, authors will be held accountable for the accuracy of all information presented as well as for the contributions. **Authors are responsible for the content of the paper they submit.** Therefore, submitting a paper with **incorrect content** (including inaccurate or hallucinated references) is grounds for rejection. [IEEE requires](#) authors to

disclose any use of AI-generated text in the paper's Acknowledgements section. The sections of the paper that contain AI-generated text must include a citation for the AI system used to generate the text.

Inclusive Description of Research Contributions

Please consider making your research contribution description more inclusive. For example, consider using examples that are ethnicity/culture-rich, consider engaging users from diverse backgrounds if your research involves a survey, etc. Make your best efforts to ensure the paper is accessible to visually impaired or color-blind readers.

Travel Arrangements

We welcome all geographically diverse sets of attendees to the conference. The authors of accepted papers must present them in person at the conference. If a visa is required for travel, the program committee will provide an invitation letter to the registered attendees. We strongly encourage the attendees to immediately apply for a visa in their country, as the process can take a significant amount of time. We are very excited to meet all of you in Seattle, USA.

IPDPS 2026 IMPORTANT DATES

- Abstract submissions: **October 1, 2026**
- Full manuscript submissions (double-anonymous): **October 8, 2026 - FIRM DEADLINE**
- Author response/rebuttal to reviews: **November 30 – December 3, 2026**
- First round decisions: **December 18, 2026**
- Revised submissions due: **January 18, 2027**
- Final decisions: **February 2, 2027**
- Camera-ready versions due: **February 20, 2027**

IPDPS 2026 PROGRAM CO-CHAIRS

- Suren Byna, The Ohio State University, USA
- Sivan Toledo, Tel-Aviv University, Israel

PROGRAM TRACK CO-CHAIRS

Algorithms

- Grey Ballard, Wake Forest University, USA
- Sanjukta Bhowmick, University of North Texas, USA

Applications

- Kamesh Madduri, Pennsylvania State University, USA
- Meifeng Lin, Brookhaven National Laboratory, USA

Architectures

- Adrian Jackson, University of Edinburgh, UK
- Angeliki Kritikakou, Inria, Rennes, France

Machine Learning and Artificial Intelligence

- Ramki Kannan, Oak Ridge National Laboratory, USA
- Nikoli Dryden, Lawrence Livermore National Laboratory, USA

Measurement, Performance and Experiments

- Nathan Tallent, Pacific Northwest National Laboratory, USA
- Rong Ge, Clemson University, USA

Programming Models, Compilers, and Runtime Systems

- Lena Oden, FernUniversität in Hagen, Germany
- Gocken Kestor, Barcelona Supercomputing Center, Spain

System Software

- Shadi Ibrahim, Inria, Rennes, France
- Jean Luca Bez, Lawrence Berkeley National Laboratory, USA

PROGRAM COMMITTEE MEMBERS

To Be Updated