

The
INTERNATIONAL CONFERENCE for **NOV 15-20**
HIGH PERFORMANCE COMPUTING
CHICAGO, IL NETWORKING, STORAGE, & ANALYSIS

SPONSORED BY  IEEE COMPUTER SOCIETY |  TCHPC |  Association for Computing Machinery |  sigahpc



SC26

Chicago, IL | **hpc**
unites.

SC26 Tech Papers Webinar

27 March 2025 • Zoom



Trilce Estrada
Associate Professor
University of New Mexico



Hatem Ltaief
Principal Research Scientist
KAUST

Paper Submissions

1 MAR 2026
Submissions Open

1 APR 2026
Abstract Submissions Close

8 APR 2026
Paper Submissions Close
(No Extensions)

24 APR 2026
AD Appendix Due
(Mandatory)

8–11 JUN 2026
Review/Rebuttal Period

1 JUL 2026
Notifications Sent

14 JUL 2026
AD Appendix & Badge
Application *(Optional)*

7 AUG 2026
Revised AD Due & AE *(if
applicable)* Due

25 AUG 2026
Artifact Freeze Deadline

28 AUG 2026
Final Paper Due

- Algorithms



Guillaume Pallez
INRIA



Giulia Guidi
Cornell University

The development, evaluation, and optimization of scalable, general-purpose, high performance algorithms.

- Applications



Tobias Weinzierl
Durham University



Alba Cristina Magalhaes Alves de Melo
University of Brasilia

The development and enhancement of algorithms, parallel implementations, models, software and problem solving environments for specific applications that require high performance resources.

- Architecture & Networks



Adrian Jackson
University of Edinburgh



Anthony Danalis
Advanced Micro Devices

All aspects of high performance hardware including the optimization and evaluation of processors and networks.

- Data analytics, visualization, storage



Ali R. Butt
Virginia Tech



Jay Lofstead
Sandia National Laboratories

All aspects of data analytics, visualization, storage, and storage I/O related to HPC systems, Submissions on work done at scale are highly favored. Further, submissions having a component focusing on the “Art of HPC” are appreciated.

- HPC for machine learning



Rio Yokota
Institute of Science Tokyo



Hari Subramoni
The Ohio State University

The development and enhancement of algorithms, systems, and software for scalable machine learning utilizing high performance computing technology. This area is primarily addressing the use of HPC to improve ML rather than the use of ML to improve any technology covered by other areas.

- Performance measurement, modeling, & tools



Tanzima Z. Islam
Texas State University



Ivona Brandic
TU Wien

Novel methods and tools for measuring, evaluating, and/or analyzing performance for large-scale systems.

- Post-Moore & Quantum Computing




Laura Schulz
Argonne National Laboratory



Nicholas Malaya
Advanced Micro Devices

Technologies that continue HPC performance scaling beyond the limits of Moore's law, including system architecture, use/integration of emerging hardware in future (HPC/AI) systems, programming frameworks, system software, and applications.

- Programming Frameworks



Bin Ren

College of William & Mary



Miwako Tsuji

University of Tsukuba / RIKEN

Compilers, programming languages, libraries, programming models, and runtime systems that enable management of hardware resources and support parallel programming for large-scale systems.

- State of the practice



Patrick Bridges
University of New Mexico



Costas Bekas
Citadel Securities

All aspects of the pragmatic practices of HPC, including operational IT infrastructure, services, facilities, large-scale application executions and benchmarks. Papers are expected to capture experiences and ongoing practice relating to modern computing centers or HPC-related software.

- System software & cloud computing



Alexandru Iosup
Vrije Universiteit Amsterdam,
Delft University of Technology

Cloud and system software architecture, configuration, optimization and evaluation, support for parallel programming on large-scale systems or building blocks for next-generation HPC architectures.



Michela Becchi
North Carolina State University

- Strengthen the rigor, transparency, and reproducibility of research, authors are required to complete a paper checklist
- Help authors reflect on key aspects of high-performance computing research, including experimental methodology, robust performance evaluation, statistical validity, and the responsible communication of limitations and impact
- The checklist is entered on the paper upload submission form as a set of yes/no questions and a short justification
- Check the sample paper online submission form [here](#)

Checklist

Authors must complete this checklist to ensure their submission meets the expectations for clarity, rigor, and reproducibility in HPC research. The checklist is designed to guide preparation and review of the manuscript.

For each item:

- Indicate Yes, No, or N/A (Not Applicable) in the corresponding question.
- If No or N/A, briefly explain the reason (e.g., “not applicable to analytical studies,” “artifact under embargo,” “baseline unavailable for architecture”).
- Responses should be specific, concise, and factual.
- Include references to the relevant sections, figures, or tables in your paper (e.g., “Section 4.2,” “Table 3”).
- All questions should be addressed explicitly, even if the answer is negative.

Examples:

Have you provided quantitative evidence or benchmarks to motivate the specific problem?

yes.

Explain: The objective is defined in Section 1.1 It emphasizes performance optimization for exascale climate simulations.

Have you identified and summarized the most relevant and recent prior work?

yes.

Explain: Section 2 reviews six state-of-the-art distributed memory solvers published between 2022–2025

Have you discussed how conclusions might change under different problem scales, architectures, or workloads?

n/a

Explain: The work focuses exclusively on shared-memory systems. Heterogeneous and distributed architectures are out of scope.

Introducing a paper checklist (3)

1. 3. 4. 5.

2. H

6.- Limitations

Have you identified the main limitations or assumptions that constrain the proposed approach?*

yes
 no
 n/a

Have you discussed how conclusions might change under different problem scales, architectures, or workloads?*

yes
 no
 n/a

0 words

Explain*

How are they addressed?*

n/a

0 words

Explain*

- Simple
- Effective
- Best practices
- Rigor

Introducing a similarity score

Submissions **My Expertise** My Conflicts Bidding Awards

My To-Do List hide
You have 1 pending task. It can be viewed [here](#).

My Expertise · prepare for assignment process

Submissions My Expertise My Conflicts **Bidding** Awards

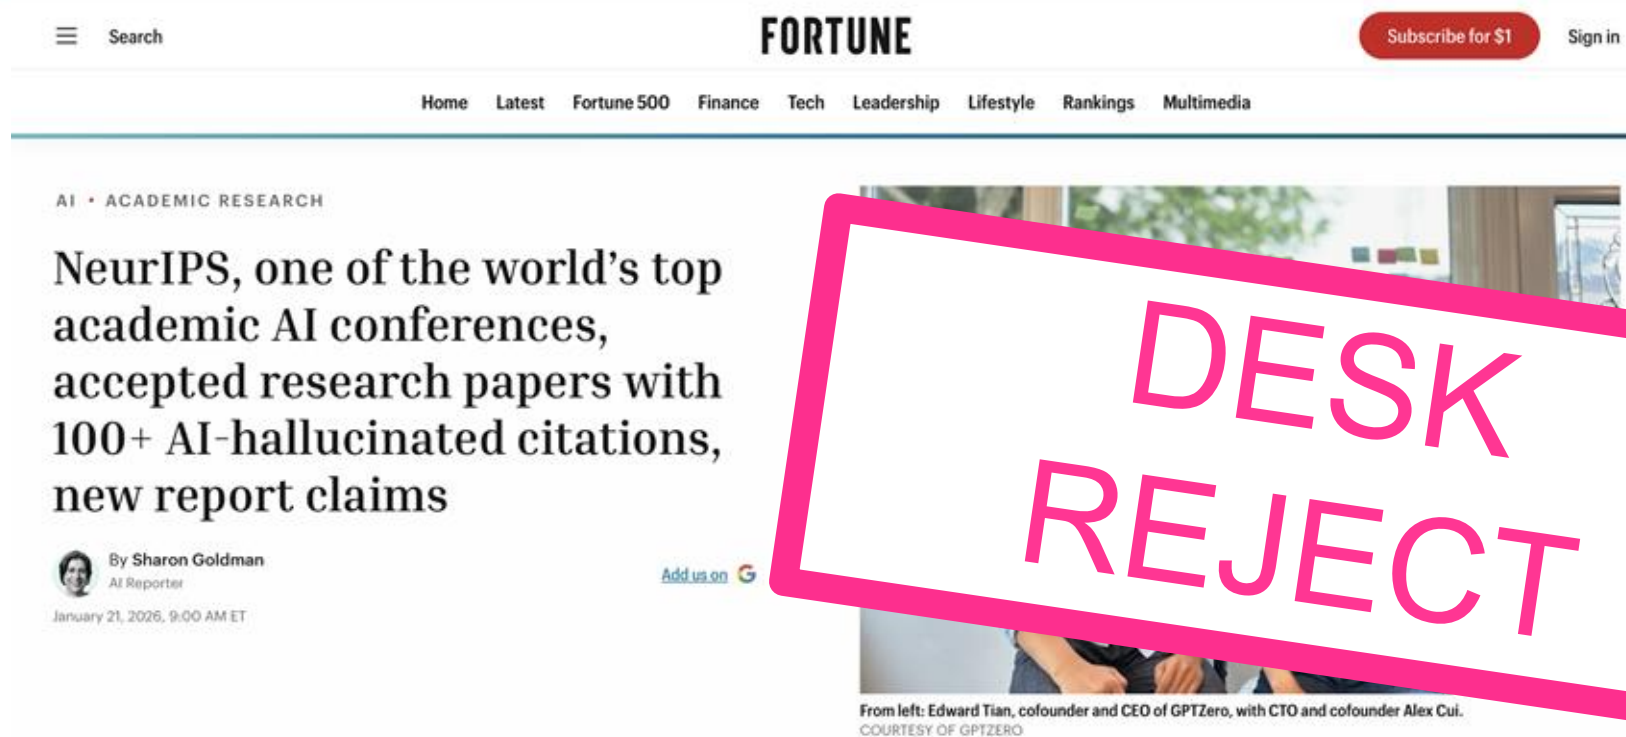
My To-Do List hide
You have 1 pending task. It can be viewed [here](#).

Bidding · request submissions I'd like to review

Indicate the submissions that you would prefer to be assigned by toggling to 'yes' in the relevant row of the 'Request?' column. Type 'f' to toggle full-screen mode.

ID	DETAILS	TITLE	TRACK / OTHERS REQ	CONFLICT?	TOPICS MATCHES - M	MATC	FILESIM - SI	REQUEST?
pap232s1	details	Dynamic Load Balancing for the Affinity-Aware Parallel Programming Model LAIK	pro	flag conflict	Programming Environr	1		
pap107s1	details	A Scheduler System for Interactive Visualization on HPC Resources	pro appl	flag conflict		0	3	no · set to yes
pap137s1	details	Runtime Context-Guided Multi-Versioning of Compute Kernels	pro algor_meth_perf	flag conflict		0	unknown	no · set to yes
pap143s1	details	Advanced Loop Kernels Optimizations within Kokkos	pro	flag conflict		0	unknown	no · set to yes
pap226s1	details	Predicting CUDA Kernel Execution Time on Edge and Server GPUs: Analytical Model vs ML Algorithr	pro algor_meth_perf	flag conflict		0	unknown	no · set to yes

FileSim - similarity score between this submission and my representative works



- Make sure to provide DOIs for all citations (it applies for most cases)
- Help us maintain research integrity
- AI tools comes with great responsibility, don't be fooled!

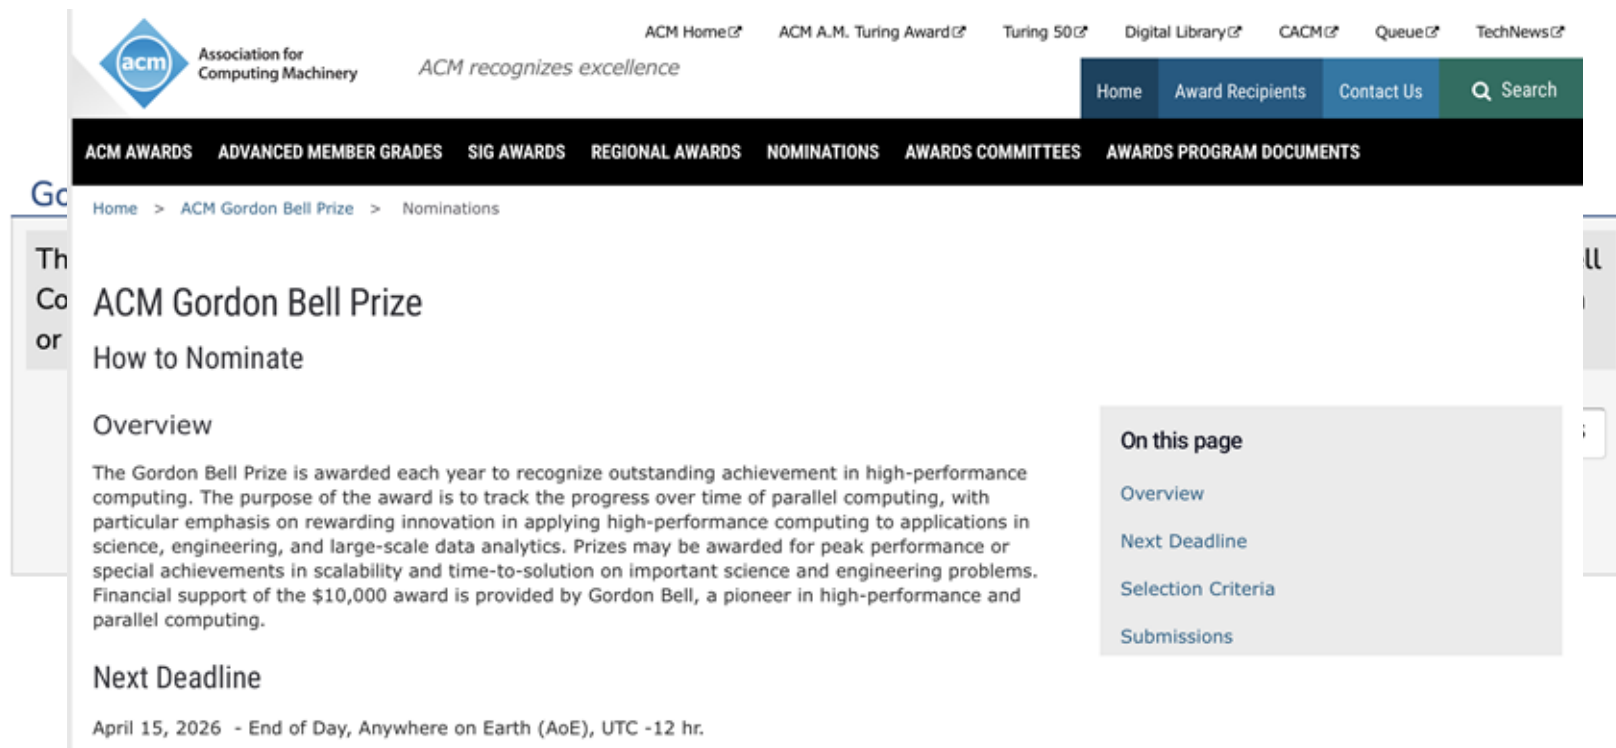
- Write your paper so as not to disclose your identity or the identities of your co-authors
- Do not use your name or your co-authors' names, affiliations, funding sources, or acknowledgments in the heading or body of the document
- Do not eliminate self-references to your published work, instead, write self-references in the third person
- Recall that the goal and spirit of double-blind review is to create uncertainty about authorship
- Declare conflicts of interest you may have with program committee members

- SC permits the use of AI tools (e.g., ChatGPT, Grammarly, or other AI assistants) to help improve your writing
- The authors must verify whether the content is accurate before submission
- AI tools are not eligible for authorship
- **Authors are completely accountable for the contents of their papers, including content generated by AI tools that could be interpreted as plagiarism or scientific misconduct (e.g., fabrication of facts)**



- Expect to uphold the highest standards of confidentiality and professional integrity
- Recruited for their expertise and judgment when assessing paper submissions
- Usage of AI tools to review papers is **strictly prohibited** as it may compromise:
 - confidentiality
 - intellectual property
 - fairness of the review process
- Provide constructive feedbacks
- We appreciate their valuable time, efforts and community services

- Submissions for the ACM Gordon Bell Prize must be made exclusively through [the official nomination portal](#) and cannot be submitted in parallel as a regular SC26 Technical Paper



The screenshot shows the ACM website's nomination portal for the Gordon Bell Prize. The header includes the ACM logo and navigation links for Home, Award Recipients, Contact Us, and Search. A secondary navigation bar lists various award categories. The main content area is titled 'ACM Gordon Bell Prize' and 'How to Nominate'. It features an 'Overview' section with a paragraph describing the prize's purpose and a 'Next Deadline' section stating 'April 15, 2026 - End of Day, Anywhere on Earth (AoE), UTC -12 hr.'. A sidebar on the right contains a 'On this page' section with links to Overview, Next Deadline, Selection Criteria, and Submissions.

- Reproducible science is essential
- Artifact Description (AD) is mandatory, Artifact Evaluation (AE) is optional
- AD/AE Appendices are integrated into the review process
- The assigned badges will be per **ACM Reproducibility Standard**.



**ARTIFACTS
AVAILABLE**



**ARTIFACTS
EVALUATED-FUNCTIONAL**



**RESULTS
REPRODUCED**



Jean Luca Bez
LBNL



Rocío Carratalá-Sáez
Universitat de València

DATES & DEADLINES

24 APR 2026: Submission of AD Appendix – (Mandatory)

14 JUL 2026: Submission of AE Appendix + Applications for Badges – (Optional)

24 JUL 2026: Feedback from AD/AE Committee on completeness of AD and AE (if applicable)

7 AUG 2026: Submission of Revised AD Appendix and Revised AE Appendix (if applicable)

25 AUG 2026: Artifact freeze, authors assign a DOI to their Artifact (Artifact must not be changed after this time)

25 SEP 2026: Artifact badge decision

Key for a good AD
Prepare it as if a student
should read and understand it

- Archived in the [ACM Digital Library](#) and [IEEE Xplore](#)
- Full papers are accessible free of charge through ACM OpenTOC, with support from SIGHPC
- Proceedings first released to attendees on the first day of the conference.
- Two volumes:
 - Tech Papers: tech papers presented at SC
 - Workshops: papers presented in SC workshops with proceedings.
- **IEEE format for all proceedings this year**



Swann Perarnau
ANL



François Tessier
INRIA



Hermes Senger
USP

Tech Papers:
28 August 2026: Final version due
No extensions

- At least one author must register for the Technical Program in order to attend the SC Conference and present the paper
- Paper presenters are expected to be in-person at the conference
- For an accepted paper to be included in the proceedings, the author has to present the paper at the conference in person. Otherwise, the paper will be removed from the proceedings
- Need a visa letter? Or have further questions? Get in touch as soon as paper notifications are released

- Before submission deadline ends
 - 10 pages not including references but it does include any non-AD/AE appendices: do not alter the template (**desk reject**)
 - Authors to ensure the paper checklist is completed and disclose COI
 - Reviewers to upload their representative works on LL
- Before review ends
 - Reviewers to self-report confidence level BEFORE reviewing
- After review ends
 - No changes in the author list allowed after the release of reviews
- After paper final notification
 - 11 pages to include the reviewers' modifications
 - Encouraged to submit camera-ready version early
 - Make sure at least one author can answer modification requests from Proceedings in September.



- We are all in this: chairs, reviewers, and authors
- We are human, we all make mistakes
- We ask your help to improve the quality of papers at SC'xy
- Our motto
 - Excellence. Responsible. Integrity. Transparency.

**Now, back to work, send us
your great research papers!**