



**Association for
Computing Machinery**

Advancing Computing as a Science & Profession

October 12–18, 2025
Singapore, Singapore



PROPL '25

Proceedings of the 2nd ACM SIGPLAN International Workshop on

Programming for the Planet

Edited by:

Anil Madhavapeddy, Dominic Orchard, and KC Sivaramakrishnan

Sponsored by:

ACM SIGPLAN, ACM SIGAda

Co-located with:

ICFP/SPLASH '25

Association for Computing Machinery, Inc.
1601 Broadway, 10th Floor
New York, NY 10019-7434
USA

Copyright © 2025 by the Association for Computing Machinery, Inc (ACM). Permission to make digital or hard copies of portions of this work for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted.

To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permission to republish from: Publications Dept. ACM, Inc.
Fax +1-212-869-0481 or E-mail permissions@acm.org.

For other copying of articles that carry a code at the bottom of the first or last page, copying is permitted provided that the per-copy fee indicated in the code is paid through the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923, USA.

ACM ISBN: 979-8-4007-2161-8

Cover photo:

Title: “Marina Bay and Merlion”,

Copyright 2025 Singapore Tourism Board, licensed to ICFP/SPLASH 2025

Production: Conference Publishing Consulting
D-94034 Passau, Germany, info@conference-publishing.com

Welcome from the Chairs

Welcome to the 2nd ACM SIGPLAN International Workshop on Programming for the Planet (PROPL 25), colocated with ICFP/SPLASH 2025 and held in Singapore on October 13th 2025! This workshop is dedicated to bridging the use of principled computer science towards positive climate and biodiversity actions. This year's edition follows a successful 1st workshop held at POPL 2024 in London, which had over 100 participants attend at the Institute of Engineering and Technology in London, UK.

Goals and Scope

There are simultaneous interlinked crises across the planet due to human actions: climate change, biodiversity loss, and desertification. Addressing these challenges requires, amongst other things, a global understanding of the present state of affairs and the effectiveness of our adaptations and mitigations, leveraging both data and computation.

However, programming the computer systems required to effectively ingest, clean, collate, process, explore, archive, and derive policy decisions from the planetary data we are collecting is difficult and leads to artefacts presently not usable by non-CS-experts, not reliable enough for scientific and political decision making, and not widely and openly available to all interested parties. Concurrently, domains where computational techniques are already central (e.g., climate modelling) are facing diminishing returns from current hardware trends and software techniques.

PROPL explores how to close the gap between state-of-the-art programming methods being developed in academia and the use of programming in climate analysis, modelling, forecasting, policy, and diplomacy. The aim is to build bridges to the current practices used in the scientific community.

Review Process

This year, we wanted to widen the accessibility of the conference to non-computer scientists, and so had three different modes of submissions:

- Provocations: (any length), short position pieces proposing and outlining a problem, application area, challenge, or capacity gap, that might be addressable by members of the community. We especially welcomed such contributions from domain experts outside computer science, and accepted submissions via an online form.

- Short papers: (up to 5 pages, excluding bibliography and appendices), addressing a topic within the scope of the workshop. We took a generous view on paper styles, given the relative youth of this workshop, so a problem statement, application or tool paper, a note on research outcomes, identification of a capacity gap or research topic are all welcome submissions. Review was conducted single-blind.
- Talk proposal: An abstract of a talk aligned with the topics of the workshop. This could include reporting on existing work, a demo, open problems, work in progress, or new ideas and speculation. Multiple talk proposals may be combined into panel discussions, depending on the submitted topics.

We merged the provocations that were suitable to be talks into the review process. At least three members of the program committee reviewed each paper and talk, and we finally accepted six papers and nine talk proposals for the workshop. There was no PC meeting, as all deliberations concluded using HotCRP discussions.

Papers

The papers in this year's workshop cover a diverse range of topics relating to programming for the planet:

- **Towards Modelling and Verification of Coupler Behaviour in Climate Models** by Chinmayi Baramashetru and Dominic Orchard explores lightweight formal verification techniques for climate model couplers to tackle correctness issues.
- **A FAIR Case for a Live Computational Commons** by Cyrus Omar, Michael Coblenz, and Anil Madhavapeddy proposes a collaborative scientific workflow called Fairground where thousands of scientists work together within a computational commons.
- **STACD: STAC Extension with DAGs for Geospatial Data and Algorithm Management** by Saharsh Laud, Saurabh Joshi, Tarun Mangla, Abhish Jindal, and Aaditeshwar Seth proposes an extension to STAC specifications that incorporates Directed Acyclic Graph representations for geospatial data workflows.
- **Bridging Disciplinary Gaps in Climate Research Through Programming Accessibility and Interdisciplinary Collaboration** by Cristian Urlea, Ana Denisa Urlea, Wim Vanderbauwhede, Adriana Laura Voinea, and

Syed Waqar Nabi addresses barriers to entry in climate research computing, particularly for interdisciplinary researchers and those in LMICs.

- **GPU-accelerated Hydrology Algorithms for On-prem Computation** by Rahul kumar, Vatsal Jingar, Abhilash Jindal, and Aaditeshwar Seth presents a GPU-accelerated framework for critical hydrology algorithms including flow accumulation, stream order, watershed delineation, and runoff simulation.
- **Yirgacheffe: a declarative approach to geospatial data** by Michael Winston Dales, Alison Eyres, Patrick Ferris, Anil Madhavapeddy, Francesca A. Ridley, and Simon Tarr presents a declarative geospatial library for spatial algorithms that supports parallel execution and automatic resource management.

Talks

In addition to the papers contained within this proceedings, we have the following talks at PROPL on a variety of topics:

- **Programming Opportunities for the Global Biodiversity Observation Network** by Jean-Michel Lord, Jamie M. Kass, Andrew Gonzalez, Michael Winston Dales, and Anil Madhavapeddy lays out the overall state of global biodiversity observations and offers provocations for computer scientists to contribute to initiatives such as BON-in-a-Box.
- **Start Making Geospatial Foundation Models Accessible** by Robin Young analyzes systematic user experience failures that prevent domain expert adoption of geospatial foundation models and outlines a research agenda for making them more accessible.
- **Challenges in Practice: Building a Usable Library for Planetary-Scale Embeddings** by Sadiq Jaffer, Frank Feng, Robin Young, Srinivasan Keshav, and Anil Madhavapeddy presents lessons learned from developing GeoTessera, an access library designed for an embeddings-as-data approach to remote sensing via the TESSERA foundation model.
- **Precision Action Towards Climate and Health (PATCH)** by Dr. Angela Chaudhuri, Nitish Kumar Venkatesan, Prerakkumar Mukeshkumar Shah, and Sabhimanvi Dua presents a comprehensive digital platform integrating multiple surveillance streams to create actionable insights for climate-induced health threats.

- **Authoring Tools for Transparent Climate Reporting** by Roly Perera, Joe Bond, Cristina David, Andrew McNutt, and Alfonso Piscitelli addresses challenges in transparent climate reporting through improved authoring tools.
- **Spatial Programming for Environmental Monitoring** by Josh Millar, Ryan Gibb, Roy Ang, Hamed Haddadi, and Anil Madhavapeddy argues for treating space as a first-class concept in programming models for large-scale environmental monitoring systems.
- **Scaling the Urban Forest: An Integrated Framework for Managing Cities by Fusing Raster and Vector Data** by Andrés C. Zúñiga-González, Anil Madhavapeddy, and Ronita Bardhan showcases a reproducible framework for building-level description of green infrastructure at national scale.
- **What we talk about when we talk about scientific programming** by Patrick Ferris explores how the scientific method relates to scientific programming and how traditional programming techniques are falling short in meeting the requirements of scientific programmers.
- **Large Language Models for Computational Climate Science** by Jay Torry balances the potential benefits of LLMs for climate analysis pipelines against their environmental costs.

Provocations

We also received several thought-provoking provocations that, while not included as formal talks or papers, raised important challenges and opportunities for the PROPL community:

- **Multi-Language Coding Environments** by Emilio Luz-Ricca (University of Cambridge) highlighted the challenges of maintaining computational reproducibility across multiple programming languages, particularly R and Python, in ecological analysis workflows. Working on global spatial analysis of hunting pressure across mammal species, Luz-Ricca noted that R users rarely provide environment information (versions, library dependencies) while Python users commonly use conda environments. The provocation identified incompatibilities between tools like RStudio IDE or VSCode with R plugins and conda environments, calling for more lightweight solutions than Docker to enable cross-language environments while maintaining reproducibility standards.

- **Forecasting on the First Mile: Computational and UX Innovations for Small-Boat Fishers' Safety in the Arabian Sea** by Maxmillan Martin (Christ University) with co-authors Prabhath H. Kurup and Sukumarapillai Abhilash presented a comprehensive case study combining high-resolution weather forecasting with participatory design for over a million artisanal fishers in India. Their work addresses the need for offline-first applications that work without connectivity while at sea, performing delivery through low-bandwidth channels (voice alerts, symbolic graphics, community radio), and the integration of traditional environmental indicators (bird patterns, wave shapes) with scientific forecasts. Through 38 interviews and 12 workshops in coastal Kerala and Tamil Nadu, they demonstrated that 5km resolution forecasts matched fishers' lived experiences better than coarse-grid models, and that fishers preferred nuanced weather information over categorical "go/don't go" advisories. The provocation emphasized how technical modeling must be combined with human-centered design, using caching tools, compressed data formats, and multimodal delivery, to create truly usable last-mile climate services.
- **Managing Natural Catastrophe Risk Needs a Systems Mindset** by Min-Si Wang (Ecomonitor) argued for reframing insurance and resilience financing as co-designed governance mechanisms rather than standalone products, particularly for Southeast Asian aquaculture farms, rice fields, and coastal communities facing typhoons and floods. The provocation critiqued current approaches that focus on individual behavior (defaults, fraud, affordability) while missing systemic fragilities across the financing stack. Wang highlighted how siloed data, static hazard curves, and black-box AI models can amplify systemic shocks when they fail to account for feedback loops between policy incentives, infrastructure, and capital flows. The provocation called for machine learning-driven geospatial analysis with rapid hazard updates, stress testing of risk cover solutions through simulation, and collaborative design between governments, insurers, and local communities; treating insurance as a tool for distributed climate governance rather than merely financial protection.
- **Improving Conservation Information in the IUCN Red List** by Ash Simkins (Cambridge Conservation Initiative) proposed leveraging AI to enhance the completeness of the IUCN Red List of Threatened Species database. The Red List contains valuable conservation information beyond extinction risk categories, including population sizes, geographic distributions, threats, and conservation actions. However, many of these fields remain incomplete,

particularly for Least Concern species, with critical information trapped in narrative text fields. The provocation highlighted challenges in extracting structured data from these narratives—where concepts like “protected area,” “national park,” and specific location names appear in various forms and languages. The proposal suggested using natural language processing and entity recognition techniques to automatically identify and code threats, conservation actions in place or needed, and pinpoint specific geographic areas within species ranges where changes, threats, or actions are present or necessary. Additionally, automated web scraping pipelines could complement IUCN data by identifying whether species face specific threats not documented in assessments, creating opportunities for developing domain-specific ontologies, knowledge graphs, and APIs that would transform the Red List into a more computationally accessible and actionable conservation resource.

These provocations represent perspectives from practitioners working at the intersection of computing and environmental challenges, and we encourage the community to consider these important problem areas in future work.

Thanks

Our thanks go to the program committee for their thorough reviews and insightful discussions during the review process. We are also grateful to all the authors for their contributions, and to the PROPL community for making this workshop possible.

Cambridge, UK
Kent, UK
Chennai, India
October 2025

Anil Madhavapeddy
Dominic Orchard
KC Sivaramakrishnan
General and PC Chairs PROPL 2025

Committee Listings

Organizing Committee

- Anil Madhavapeddy (University of Cambridge)
- Dominic Orchard (University of Kent)
- KC Sivaramakrishnan (IIT-Madras)

Program Committee

- Chinmayi Prabhu Baramashetru (University of Oslo)
- Valentin Churavy (MIT CSAIL)
- Justin Hsu (Cornell University)
- Roly Perera (University of Cambridge/University of Bristol)
- Benjamin C. Pierce (University of Pennsylvania)
- Lisa Rennels (University of California at Berkeley)
- Aaditeshwar Seth (Indian Institute Of Technology Delhi)
- Lauritz Thamsen (University of Glasgow)
- Michele Weiland (University of Edinburgh)

Contents

Frontmatter

Welcome from the Chairs	iii
PROPL 2025 Organization	ix

Papers

Towards Modelling and Verification of Coupler Behaviour in Climate Models Chinmayi Baramashetru and Dominic Orchard — <i>University of Kent, UK; University of Cambridge, UK</i>	1
A FAIR Case for a Live Computational Commons Cyrus Omar, Michael Coblenz, and Anil Madhavapeddy — <i>University of Michigan, USA; University of California at San Diego, USA; University of Cambridge, UK</i>	8
STACD: STAC Extension with DAGs for Geospatial Data and Algorithm Management Saharsh Laud, Saurabh Joshi, Tarun Mangla, Abhilash Jindal, and Aaditeshwar Seth — <i>IIT Delhi, India</i>	14
Bridging Disciplinary Gaps in Climate Research through Programming Accessibility and Interdisciplinary Collaboration Cristian Urlea, Ana Denisa Urlea, Wim Vanderbauwhede, Adriana Laura Voinea, and Syed Waqar Nabi — <i>University of Glasgow, UK; Romanian Air Traffic Services Administration, Romania</i>	24
GPU-Accelerated Hydrology Algorithms for On-Prem Computation: Flow Accumulation, Drainage Lines, Watershed Delineation, Runoff Simulation Rahul Kumar, Vatsal Jingar, Abhilash Jindal, and Aaditeshwar Seth — <i>IIT Delhi, India</i>	31
Yirgacheffe: A Declarative Approach to Geospatial Data Michael Winston Dales, Alison Eyres, Patrick Ferris, Francesca A. Ridley, Simon Tarr, and Anil Madhavapeddy — <i>University of Cambridge, UK; Newcastle University, UK; IUCN, UK</i>	47
Author Index	55