Preface

This volume contains the Proceedings of the 40th International Conference on Mathematical Foundations of Programming Semantics (MFPS XL), which was held from 19th to 21st June 2024 at the University of Oxford in the UK. The conference was colocated with the 7th International Conference on Applied Category Theory (ACT), which took place from 17th to 21st June. While the conference was primarily organized around attendance in person, support for on-line attendees was also provided, with all talks live-streamed and a few talks given remotely. We thank the local organizer, Sam Staton, and his team (Pedro Amorim, Benedict Bunting, Younesse Kaddar, Anthony Kattuman, Jacek Karwowski, Jack Liell-Cock, Paolo Perrone, Kashish Raimalani, Jessica Richards, Benjamin Rodatz, Mario Roman, Philip Saville, Razin Shaikh, Zev Shirazi, Ned Summers, Ruben Van Belle, Theo Wang, Jingjie Yang and Lara Hartley) for the smooth running of the event, and for ensuring a successful and above all enjoyable first colocation of MFPS and ACT. We also thank the MFPS Organizers (Steering Committee), in particular Michael Mislove, for their guidance and advice.

MFPS conferences are dedicated to the areas of mathematics, logic, and computer science that are related to models of computation in general, and to semantics of programming languages in particular. This is a forum where researchers in mathematics and computer science can meet and exchange ideas. Topics include, but are not limited to, the following: bio-computation, concurrent qualitative and quantitative distributed systems, process calculi, probabilistic systems, constructive mathematics, domain theory and categorical models, formal languages, formal methods, game semantics, lambda calculus, programming-language theory, quantum computation, security, topological models, logic, type systems and type theory.

MFPS 2024 continued the tradition of having an exciting group of invited speakers, both plenary and in special sessions. We enjoyed listening to Andrew Pitts, who gave a joint plenary talk with ACT, and to Robert Atkey, Philippa Gardner, Ohad Kammar and Catuscia Palamidessi, who all gave MFPS keynote talks.

MFPS and ACT held a joint special session in memory of Phil Scott, following his untimely passing in December 2023. This was organized by Rick Blute with invited contributions from Rick Blute, Samson Abramsky, Robin Cockett and Mark Lawson. From the MFPS perspective, this was an opportunity to pay our respects to someone who was close to many of us in the MFPS community, whose academic work had great influence on several MFPS areas, and who was also known to have a personal fondness for the MFPS conference series.

A further MFPS special session on the semantics of non-wellfounded and circular proofs was organised by Anupam Das and Abhishek De, with invited contributions from Farzad Jafarrahmani, Farzaneh Derakhshan and Gianluca Curzi.

For the first time, MFPS XL invited two categories of contributed submission: *regular research contributions* and *early announcements*. The latter provided an opportunity to present novel work at an early stage of development. Contributions in this category were given the same presentation time (30 minutes) as regular research contributions.

The conference received 30 submissions in total, comprising 26 regular research contributions and 4 early announcements. Submissions were received from Australia, Brazil, Canada, the Czech Republic,

Published December 15, 2024 Available Online at https://doi.org/10.46298/entics.proceedings.mfps40

Denmark, France, Germany, Italy, North Macedonia, The Netherlands, South Korea, the United Kingdom and the United States. The program committee accepted 22 submissions for presentation: 19 regular research contributions and 3 early announcements. We are very grateful to the program committee members and external reviewers for their exceptional work over an extremely tight time schedule. A list of authors and titles of the accepted early announcements is given at the end of this preface. The main content of this proceedings volume then follows: the regular research contributions, which are published in this volume in alphabetical order of the names of first authors. We hope you enjoy reading them!

Valeria de Paiva and Alex Simpson
December 2024

Program Committee

Danel Ahman	*University of Tartu*
Amal Ahmed	*Northeastern University*
Sandra Alves	*University of Porto*
Paolo Baldan	*Università di Padova*
Andrej Bauer	*University of Ljubljana*
Lars Birkedal	*Aarhus University*
Steve Brookes	*Carnegie Mellon University*
Valentina Castiglioni	*Eindhoven University of Technology*
Liron Cohen	*Ben-Gurion University*
Raphaëlle Crubillé	*IRIF*
Valeria de Paiva (co-chair)	*Topos Institute*
Masahito Hasegawa	*Kyoto University*
Justin Hsu	*Cornell University*
Patricia Johann	*Appalachian State University*
Achim Jung	*University of Birmingham*
Marie Kerjean	*CNRS, Univ. Sorbonne Paris Nord*
Robbert Krebbers	*Radboud University*
Paul Levy	*University of Birmingham*
Michael Mislove	*Tulane University*
Daniele Nantes-Sobrinho	*Imperial College London*
Alex Simpson (co-chair)	*University of Ljubljana*
Christine Tasson	*Sorbonne Université*

List of Subreviewers

Giorgio Bacci, Pablo Barenbaum, Eduardo Bonelli, Pierre Cagne, Jules Chouquet, Abhishek De, Jim de Groot, Paulo Emílio de Vilhena, Ryan Doenges, Thomas Ehrhard, Daniel Gratzer, Charles Grellois, Ariel Grunfeld, Daniel Hillerström, Kuen-Bang Hou Favonia, Dominique Larchey-Wendling, John Li, Edward Morehouse, Federico Olimpieri, Andrew Polonsky, James Riely, Nicholas Rioux, Steffen van Bakel, Daniel Ventura, Andrew Wagner.

Invited Plenary Talks

Andrew Pitts (University of Cambridge, UK) *Toposes of finitely supported M-sets* (Joint with ACT)

For at least the last 25 years several toposes of finitely supported M-sets (for various monoids M) have played an important role in my work on logic and computer science, but in apparently unconnected ways. I feel that this kind of topos is trying to get my attention. So I will use this talk to tell you what they are; some of their nice (and not so nice) properties; and a little about the contexts in which they have arisen for me, namely the mathematics of syntax involving binders and the semantics of univalent type theories. My hope is that you will also find them useful, or be able to tell me something about them that I don't know. For example, I do not know an abstract characterisation of this class of topos.

Robert Atkey (University of Strathclyde, UK) *Polynomial time and dependent types*

Dependent Type Theory is inherently computational. Constructions and proofs in Type Theory are programs that can be executed to produce results. However, all these programs and explanations are built with no regard for computational complexity within the theory, meaning that constructions cannot reason or rely on complexity bounds, and arguments that rely on complexity restrictions lie outwith the grasp of Type Theory. To address this, I will talk about combining techniques from Implicit Computational Complexity with Dependent Type Theory to attempt to bring resource consciousness to Type Theory.

Philippa Gardner (Imperial College, London, UK) *Computational symbolic execution for over-approximating and under-approximating reasoning*

A relatively recent challenge has been to develop symbolic-execution techniques and tools that are functionally compositional with simple function specifications that can be used in broader calling contexts. The technical break-through came with the introduction of separation logics for reasoning about partial mutable state, leading to compositional symbolic execution tools being developed in academia and industry. Many of these tools have been grounded on a formal foundation, but either the function specifications are validated with respect to the underlying symbolic semantics of the engine, with no meaning outside the tool, or there is a large gulf between the theory and the implementations of the tools. In this talk, I will introduce a formal compositional symbolic execution engine which creates and uses function specifications from an underlying separation logic and provides a sound theoretical foundation for, and indeed was partially inspired by, the Gillian platform. This is achieved by providing an axiomatic interface which describes the properties of the consume and produce functions for updating the symbolic state when calling function specifications, a technique used by VeriFast, Viper and Gillian but not previously characterised independently of the tool. A surprising property is that our semantics provides a common foundation for both correctness and incorrectness reasoning, with the difference in the underlying engine only amounting to the choice to use satisfiability or validity. We use this insight to extend the Gillian platform with incorrectness reasoning, developing automatic true bug-finding using incorrectness bi-abduction, which our engine incorporates by creating fixes from missing-resource errors. We have shown that the Gillian implementation of the consumer and producer functions satisfy the properties described by our axiomatic interface, and evaluate our new Gillian platform by using the Gillian instantiation to C. This instantiation is the first tool to support both correctness and incorrectness reasoning, as well as being grounded on a common formal compositional symbolic execution engine. (Authors: Andreas Lööw, Daniele Nantes Sobrinho, Sacha-Elie Ayoun, Caroline Cronjäger, Petar Maksimović and Philippa Gardner.)

Ohad Kammar (University of Edinburgh, UK) *Semantic foundations for type-driven probabilistic modelling*

The last few years have seen several breakthroughs in the semantic foundations of probabilistic and statistical modelling. Types show clear promise in organising intricate models and the inference algorithms

we use to fit them to data. I will present a type-rich and straightforward model, the quasi Borel space, and survey recent and ongoing developments in this area.

Catuscia Palamidessi (INRIA Saclay & LIX, France) *Information structures for privacy and fairness*

The increasingly pervasive use of big data and machine learning is raising various ethical issues, in particular privacy and fairness. In this talk, I will discuss some frameworks to understand and mitigate the issues, focusing on iterative methods coming from information theory and statistics. In the area of privacy protection, differential privacy (DP) and its variants are the most successful approaches to date. One of the fundamental issues of DP is how to reconcile the loss of information that it implies with the need to preserve the utility of the data. In this regard, a useful tool to recover utility is the Iterative Bayesian Update (IBU), an instance of the famous Expectation-Maximization method from Statistics. I will show that the IBU, combined with the metric version of DP, outperforms the state-of-the art, which is based on algebraic methods combined with the Randomized Response mechanism, widely adopted by the Big Tech industry (Google, Apple, Amazon, ...). Furthermore I will discuss a surprising duality between the IBU and one of the methods used to enhance metric DP, that is the Blahut-Arimoto algorithm from Rate-Distortion Theory. Finally, I will discuss the issue of biased decisions in machine learning, and will show that the IBU can be applied also in this domain to ensure a fairer treatment of disadvantaged groups.

Special Sessions

Joint ACT/MFPS Special Session in Memory of Phil Scott

The session was organized by Rick Blute and contained the following invited talks:

Rick Blute (University of Ottawa, Canada) *Proofs, types and hexagons*

Phil's book with Jim Lambek *Introduction to Higher-Order Categorical Logic* is a seminal work in category theory. It cemented in the minds of many of us working in the field the idea that the best approach to studying logical systems was to form a category whose objects are formulas and arrows are proofs of entailments. Extensions of this idea are an obsession of the categorical logic community to this day. One of the most interesting extensions involves representing formulas as multivariant functors and thus proofs become dinatural transformations, the hexagons of the title. The problem/challenge that arises is that these transformations need not compose. This has led to a great deal of interesting category theory, much of which Phil contributed to. I'll review some of these results, especially those relating to linear logic.

Samson Abramsky (University College London, UK) *Retracing GoI with Phil*

I will describe my collaboration with Phil Scott and Esfandiar Haghverdi on categorical semantics of Geometry of Interaction. I will say something about the context of this work, and where it led.

Robin Cockett (University of Calgary, UK) *Phil Scott and Occam's razor*

In Phil's book with Lambek the authors confess they have "axes to grind". One of these is "We decry over zealous application of Occam's razor." Despite this, the book deploys Occam's razor to good effect in order to describe how models of the untyped lambda calculus can be derived. The talk traces these ideas forward leading into how they influenced the speaker's work.

Mark Lawson (Heriot-Watt University) *Co-ordinatizing an MV-algebra by a Boolean inverse monoid*

Both MV-algebras (which arise from multiple valued logic) and Boolean inverse monoids (which partake in a generalization of classical Stone duality) are generalizations of Boolean algebras. The question is: how are they related? Phil Scott and I proved that every countable MV-algebra can be co-ordinatized by a suitable Boolean inverse monoid (in a sense analogous to classical work by von Neumann). I shall assume no prior exposure to MV-algebras or Boolean inverse monoids in my talk.

Special Session on the Semantics of Non-wellfounded and Circular Proof

The session was organized by Anupam Das and Abhishek De and contained the following invited talks.

Farzad Jafarrahmani (LMCRC, Huawei) *Focused orthogonality as denotations of circular and non-wellfounded proofs*

This talk investigates the question of denotational invariants of non-wellfounded and circular proofs of the linear logic with least and greatest fixpoints. While non-wellfounded and circular proof theory has made significant progress in the last twenty years, the corresponding denotational semantics is still underdeveloped. First, we explore a theory of fixpoint constructions in focused orthogonality categories and present a lifting theorem for initial algebras and final coalgebras. These constructions crucially hinge on the insight that focused orthogonality categories are relational fibrations. We then demonstrate that assuming a CPO structure on our category allows the focused orthogonality construction to provide a model for non-wellfounded proofs. Several properties of the semantics will be discussed, including its soundness, the relationship between totality (orthogonality) and validity, and the semantic content involved in translating finitary proofs to circular proofs. Finally, the talk focuses on circular proofs, aiming to leverage their regularity to define the interpretation function inductively. We argue why the usual validity condition

is too general for this purpose, while a fragment of circular proofs—strongly valid proofs—constitutes a well-behaved class for such an inductive interpretation.

Farzaneh Derakhshan (Illinois Tech) *Session-typed recursive processes and infinitary proofs*

Session types describe the communication behavior of interacting processes. Binary session types, in which each channel has two endpoints, corresponds to intuitionistic linear logic by a Curry-Howard interpretation of propositions as types, proofs as programs, and cut reduction as communication. This interpretation provides a solid foundation for reasoning about the behavior of session-typed processes. For example, termination of a process can be inferred from the cut elimination property of its underlying proof. However, as soon as we add recursive session types, we abandon this correspondence and lose our solid ground. From the programming perspective, it means that we can no longer exploit the cut elimination property to guarantee termination. In this talk, I show how we can revitalize the logical foundation for recursive binary session-typed processes by using infinitary proof systems for linear logic. We show that if we refine recursive types as least and greatest fixed points and impose a guard condition on recursive processes, we can still guarantee meaningful communication, ensuring that a program always terminates either in an empty configuration or one attempting to communicate along external channels.

Gianluca Curzi (University of Gothenburg) *Computational aspects of cyclic and non-wellfounded proofs*

Non-wellfounded proof theory studies a more permissive notion of proof seen as a possibly infinite (but finitely branching) tree structure. Cyclic proofs are special non-wellfounded proofs whose underlying tree is regular, thereby they admit a finite presentation in terms of finite (possibly cyclic) graphs. Over the years, cyclic and non-wellfounded proofs have received growing interest, especially given their capability of subsuming various forms of (co)inductive reasoning and, under a computational reading, (co)recursion mechanisms. This talk is meant to provide a bird's eye view of the state-of-the-art research around the computational analysis of cyclic and non-wellfounded proof systems, with special attention to the speaker's contribution on the topic. More specifically, the talk will examine a number of systems, spanning from those expressing complex forms of ordinal and higher-order recursion down to those implementing quickly converging recursion schemes suitable for complexity-theoretic considerations.

Early Announcements

The following *Early Announcements* were presented, with the speaker's name in bold:

Willem Heijltjes (University of Bath) *The Functional Machine Calculus III: Choice*

The Functional Machine Calculus (Heijltjes 2022) is an extension of the lambda-calculus that preserves confluent reduction and typed termination, while enabling both call-by-name and call-by-value reduction behaviour and encoding the computational effects of mutable higher-order store, input/output, and probabilistic computation. In this work the calculus is extended to capture exception handling and loop constructs.

Patricia Johann and **Pierre Cagne** (Appalachian State University) *Relational parametricity in the presence of GADTs*

Relational parametricity was first introduced by Reynolds for System F. Although System F provides a strong model for the type systems at the core of modern functional programming languages, it lacks features of the daily programming practice such as complex data types. In order to reason parametrically on such objects, Reynolds' seminal ideas need to be generalized to extensions of System F. Here, we explore such a generalization for the extension of System F by Generalized Algebraic Data Types (GADTs) as found in Haskell. Although GADTs generalize Algebraic Data Types (ADTs), i.e., simple recursive types such as lists, trees, etc., we show that it is not enough to naively extend the parametric treatment of these recursive types to tackle GADTs. We propose a tentative workaround for this issue, borrowing ideas from the categorical semantics of GADTs known as "functorial completion". We discuss some applications, as well as some limitations, of this solution.

Jack Liell-Cock and Sam Staton (University of Oxford) *Compositional models for imprecise probability*

"Imprecise probability" is concerned with uncertainty about which probability distributions to use. It has applications in robust statistics and elsewhere. Imprecise probability can be modelled in various ways, including by convex sets of probability distributions.

We look at programming language models for imprecise probability. Our desiderata are that we would like our model to support all kinds of composition, categorical and monoidal, in other words guided by dataflow diagrams. Another equivalent perspective is that we would like a model of synthetic probability in the sense of Markov categories.

There is already a fairly popular monad-based approach to imprecise probability, but it is arguably not fully compositional because the monad involved is not commutative, which means that we do not have a proper monoidal structure. In this work we provide a new fully compositional account. The key idea is to name the non-deterministic choices. To manage the renamings and disjointness of names, we use graded monads. We show that the resulting compositional model is maximal. We relate with the earlier monad approach, showing that we obtain tighter bounds on the uncertainty.

This is a submission in the "early announcement" category. Imprecise probability and the questions of combining non-determinism and probability have arisen many times at MFPS over the years, and we are keen to discuss our seemingly new approach with the MFPS community.