

Mark Shields (PhD) Resume

[<https://bit.ly/3rhqPB6>, last updated Oct-4-2023]

Software Engineer, [Modular](#).

mshields822@gmail.com

<https://www.linkedin.com/in/mshields822>

Seattle, WA, USA



- Domains** Compilers & languages (ML, FP, DSLs, runtimes), maps (routing, conflation, real-time processing), distributed dataflow (cf Spark).
- Technology** Languages and type systems theory and practice, deep-learning frameworks/compilation/runtime, distributed systems for big data, C++, Python, Java, C#, OCaml, Haskell.
- People** Team building and management, area tech lead / architect, mentoring, excellent verbal and written communication skills, strong sense of hustle.
- Work status** Australian, Greencard
- Hobbies** Backcountry hiking, trail running, house renovation, electronics.

Employment

- Aug 2022-Present** **Software Engineer, [Modular](#)**. I work in the intersection of the inference engine runtime, graph compiler, and Mojo language to ensure the Modular stack can import, compile and execute customer TensorFlow/Pytorch/ONNX models with state-of-the-art performance across both CPU and GPU. C++, MLIR, CUDA, DL, memory planning/optimization/management, CPU concurrency, CPU/GPU concurrency.
- Jul 2021-Jul 2022** **ML Compiler Engineer, [OctoML](#)**. I led the 'CORE' team responsible for development and maintenance of the [TVM](#) front-end. This consists of the AST and passes around [Relay](#) (resembling a functional language like OCaml), and some of the machinery responsible for lowering from Relay to TIR (resembling a [Hallide](#)-like imperative tensor processing language). I've re-implemented [device planning](#) for heterogeneous backend support. I've made extensive refactors to prepare for incremental lowering, with a view to tackling global layout and memory scope optimizations by flowing constraints between lowered kernels and the yet-to-be-lowered graph. I've since started the [Collage](#) project to bring search and flexibility to partitioning and fusion, with particular emphasis on allowing global tuning to mix and match between libraries, 3rd party toolchains and TVM itself for partition/kernel implementations. C++, Python, Deep Learning, optimization, codegen. [Github](#).
- Aug 2019-Jun 2021** **Area Tech Lead (Senior Staff Software Engineer), [Google/Maps/Routing](#)**. I designed, prototyped, and led the productionization of a new routing coordination system to improve our support for 'customized' routing. The system is responsible for rewriting routing queries into an internal query plan, gathering candidate routes from multiple route generation backends, gathering annotations required to score the routes from multiple ML models,

scoring, and final ranking. I also guided the design and implementation of a new service for enterprise customers to request custom routes for ride-share and delivery. I oversaw development across route generation, ranking and the various feedback loops by which we assess and improve routing quality. C++, microservices, routing algorithms.

**Dec 2016-
Jul 2019**

Area Tech Lead (Staff Software Engineer), Google/Maps/Session Analytics. Designed and implemented a new telemetry, analysis and metrics system for Google Maps for Mobile. Helped feed analysis output back into product quality improvements for route selection. This system now actively detects and repairs maps issues which cause 'deviations' and other bad user outcomes. Managed a team which replaced heuristic guidance systems with ML models. C++, BigData, HMMs.

**Jun 2016-
Nov 2016**

Sabbatical. Took some time off to catch up on statistical programming, Bayesian statistics, ML, and work on my house.

**Jul 2015-
May 2016**

IC (Staff Software Engineer), Google/Cloud/Dataflow. Helped launch [Google Cloud Dataflow](#), with focus on the streaming runtime. Established benchmark suites. Made many performance, concurrency, correctness and semantics improvements. Onboarded EAP customers. Helped transition SDK to an [Apache 'Beam' project](#). Java, C++, Google Compute Engine. A little Spark, Flink and Scala for cross-comparing benchmarks.

**Apr 2011-
Jul 2015**

Tech Lead/Manager (Senior Software Engineer), Google/Maps/Ground Truth. Pioneered model-based data repair to improve geocoding and business listing locations. Built the pipeline which consumes [auto-transcribed street numbers](#) into the base map database. Built and managed a team of six engineers. Java, BigData, custom algorithms.

**Aug 2010-
Mar 2011**

IC (Senior Developer), Microsoft/DevDiv. Worked on the new [Chakra](#) JavaScript engine in Internet Explorer 9. Wrote the new regular-expression engine. C++.

**Jun 2007-
Aug 2010**

IC (Senior Developer), Microsoft/SQL/Volta. Built an optimizing compiler to transpile .Net IL to JavaScript, and implemented the corresponding runtime. Implemented updatable views on SQL backend via C# LINQ. C#, JavaScript.

**Jan 2006-
May 2007**

Program Manager, Microsoft/SQL/XML. Prototyped a data modeling language combining relational and XSD concepts. OCaml.

**Jul 2002-
Dec 2006**

Applied research scientist, Galois Inc. Startup. Developed new consulting business with various defense clients. Multi-level security system design for verifiability. Built a certifying compiler and its runtime for a domain specific language for block-cipher cryptographic algorithms on custom verified hardware. Haskell, OCaml, C.

**Apr 2002-
May 2002**

Research Fellow, Department of Computer Science, Melbourne University, Australia. Research in static analysis.

**Feb 2001-
Apr 2002**

Post-doctoral researcher, Microsoft Research, Cambridge, United Kingdom. Research in advanced type systems for functional programming languages.

**Feb 1996-
Feb 2001**

Student. PhD in Computing Science, Oregon Graduate Institute, University of Glasgow, Sydney Institute of Technology.

1994-1995

Student. Bachelor of Science Honors Student in Computing Science, University of Melbourne, Australia.

- Sep 1993-
Mar 1995** **Software Engineer, Australian Bionic Ear and Hearing Research Institute, Melbourne, Australia.** Programming support for ear implant customization. C++.
- Apr 1992-
Aug 1993** **Software Engineer, Systematix Pty Ltd, Melbourne, Australia.** Startup. Cash delivery optimization software. C++.
- Aug 1991-
Apr 1992** **Software Engineer, Co-Cam Computer Group, Melbourne, Australia.** Real-time financial information systems. C.
- Mar 1989-
Aug 1991** **Software Engineer, Systematix Pty Ltd, Melbourne, Australia.** Startup. Business programming.
- Jun 1988-
Mar 1989** **Programmer, Arthur Andersen & Co, Melbourne, Australia.** Business programming.
- 1987-1990** **Student.** Bachelor of Science in Computer Science, Monash University, Australia.

Papers

- 2007** [Practical type inference for arbitrary-rank types](#). Simon Peyton Jones, Dimitrios Vytiniotis, Stephanie Weirich, and Mark Shields. JFP 17(1)
- 2006** [A verifying core for a cryptographic language compiler](#). Lee Pike, Mark Shields, John Matthews. In ACL2'06.
- 2006** [A language for symmetric-key cryptographic algorithms and its efficient implementation](#). Mark Shields. Galois Connections Technical Report.
- 2002** [Lexically-scoped type variables](#). Simon Peyton Jones and Mark Shields. Unpublished.
- 2002** [A compiler writer's guide to C#](#). Mark Shields. Lecture notes.
- 2002** [First-class modules for Haskell](#). Mark Shields and Simon Peyton Jones. In FOOL 9.
- 2001** [Object-Oriented style overloading for Haskell](#). Mark Shields and Simon Peyton Jones. In BABEL'01.
- 2001** [Static types for dynamic documents](#). Mark Shields. PhD thesis.
- 2001** [Type-indexed rows](#). Mark Shields and Erik Meijer. In POPL'01.
- 2000** [XMLambda: A functional programming language for constructing and manipulating XML documents](#). Erik Meijer and Mark Shields. Unpublished.
- 2000** [Implicit parameters: Dynamic scoping with static types](#). Jeffrey Lewis, Mark Shields, Erik Meijer and John Launchbury. In POPL'00.
- 1998** [Dynamic typing as staged type inference](#). Mark Shields, Tim Sheard and Simon Peyton Jones. In POPL'98.
- 1998** [Bridging the gulf: A common intermediate language for ML and Haskell](#). Simon Peyton Jones, John Launchbury, Mark Shields and Andrew Tolmach. In POPL'98. I've personally retracted this paper, see the [corrigendum](#).

Patents

- 8051426** **Co-routines native to a virtual execution environment.** Henricus Johannes Maria Meijer, Brian C. Beckman, Mark B. Shields.
- 8209674** **Tier splitting support for distributed execution environments.** Henricus Johannes Maria Meijer, Brian C. Beckman, Christopher W. Brumme, Mark B. Shields, Wei Zhu.
- 8464280** **Execution context control.** John Dyer, Henricus Johannes Maria Meijer, Mark Shields, Jeffrey van Gogh, Danny van Velzen, Brian Beckman, Harish Kantamneni.
- 8978007** **Static metadata in dynamic programs.** Lucas J. Hoban, Mark B. Shields, Steven E. Lucco, Charles P. Jazdzewski, Anders Hejlsberg.

Education

- Feb 2001** Ph.D. in Computing Science, Oregon Graduate Institute, USA. Functional programming, type systems, DSLs.
- Mar 1996** B.Sc. (hons) in Computing Science, University of Melbourne, Australia.
- Mar 1991** B.Sc. in Computing Science (with 1st year Engineering and Music), Monash University, Australia.