

# George Karpenkov

SOFTWARE ENGINEER · COMPILERS FOR ML ON GPU

✉ [george@metaworld.me](mailto:george@metaworld.me) | 🏠 <https://www.metaworld.me> | 🗣️ cheshire | 🌐 georgekarpenkov

## Skills

---

**Areas of Interest** Compilers, Optimizations, Systems Programming, Machine Learning, GPU, Convex Optimization

**Technology** LLVM, XLA, MLIR, Triton

**Languages** C/C++, Python, JAVA

## Experience

---

### Google

*Zurich, Switzerland*

STAFF SOFTWARE ENGINEER

*2019 - Now*

- Leading XLA:GPU performance work. Created a GPU performance team from zero to nine in a new location.
- Led the project to replace direct-to-LLVM XLA:GPU emitters with Triton, resulting in 3-4x speedups on LLVM inference benchmarks
- Revamped development methodology for XLA:GPU, heavily investing into benchmarking infrastructure and a metric-driven approach
- XLA:GPU contact point with NVIDIA, guiding NVIDIA contributions into XLA:GPU.
- Led XLA:GPU integration into TensorFlow using `tf.function(jit_compile=True)`, resulting in >2x speedups on TF workloads.
- Revamped documentation on [openxla.org](https://openxla.org)

### Apple

*Cupertino, California*

SOFTWARE ENGINEER

*2017 - 2019*

- Worked on Clang Static Analyzer (Build->Analyze in XCode) and libFuzzer
- Integrated libFuzzer into Clang toolchain, making `-fsanitize=fuzzer` flag possible.
- Integrated libFuzzer into Swift, enabling easier fuzzing of Swift projects out of the box.
- Rewrote the exploration order of Clang Static Analyzer from naive BFS to coverage-based priority queue, getting 20% increase in number of bugs found
- Mentored the GSoC project for cross-checking static analyzer reports using Z3, getting a 10% increase in precision

## Talks

---

### Triton Conference

*Microsoft Silicon Valley Campus*

TRITON USAGE IN XLA:GPU EMITTERS

*2023*

- Talk: <http://youtu.be/UIvM0eNdrQ>

### OpenXLA Summit

*Google Sunnyvale Campus*

XLA:GPU ARCHITECTURE

*2023*

- Talk: <http://youtu.be/xGxq0pueAXE>

### TensorFlow Summit

*Google Sunnyvale Campus*

USING XLA:GPU COMPILATION IN TENSORFLOW

*2023*

- Talk: <http://youtu.be/cPAD9vLKE0c>
- >30K Views and >300 Likes

### WWDC 2018

*San Jose Convention Centre*

WHAT'S NEW IN STATIC ANALYZER

*2018*

- Presenting advances in Clang Static Analyzer

### LLVM Developer Conference

*San Jose Convention Centre*

IMPROVED STATIC ANALYZER

*2018*

- Talk: <http://youtu.be/4n31-ZcDJNY>

## Education

---

### Verimag Laboratory

*Grenoble, France*

PHD IN COMPUTER SCIENCE

*2013 - 2017*

- Developing methods for automated software verification and bug finding (formal methods) using convex programming (policy iteration approach).
- Created a new approach for invariant finding, which took the first place in the “Loops” category of the International Competition on Software Verification
- Worked on **CPAChecker** ([cpachecker.sosy-lab.org](http://cpachecker.sosy-lab.org)), a collaboratively developed open-source tool for automated software verification and bug-finding
- Created **JavaSMT** ([github.com/sosy-lab/java-smt](https://github.com/sosy-lab/java-smt)) library, a unified library for using SMT (satisfiability modulo theories) solvers from JAVA
- Using convex programming for software verification. Thesis: [metaworld.me/these.pdf](http://metaworld.me/these.pdf)

### The University of Sydney

*Sydney, Australia*

BACHELOR OF IT, PHYSICS AND MATHEMATICS

*2008 - 2013*

## Honors & Awards

---

- 2012 **1st Place in Australia and New Zealand in ACM International Collegiate Programming Contest,** *Sydney, Australia*  
ICPC 2012

## Publications

---

- 2016 **JavaSMT: A Unified Interface for SMT Solvers in Java**, VSTTE  
2016 **Formula Slicing: Inductive Invariants from Preconditions**, HVC  
2016 **Program Analysis with Local Policy Iteration**, VMCAI