

Lechen Yu

Education

- 2017 - 2024 **Doctor of Philosophy (will graduate in August, 2024)**, *Georgia Tech*, Atlanta, GA, USA.
Research Advisor: Prof. Vivek Sarkar, College of Computing, Georgia Tech
Major: Computer Science (focusing on high performance computing), GPA: 3.75/4.00
PhD Thesis: Dynamic Analysis of Data Inconsistencies and Data Races in OpenMP Programs
- 2015 - 2017 **Master of Science**, *Rice University*, Houston, TX, USA.
Research Advisor: Prof. Vivek Sarkar, Department of Computer Science, Rice University
Major: Computer Science, GPA: 3.92/4.00
Master's Thesis: Data Race Detection for Event-Driven Parallel Runtime Systems
- 2012 - 2015 **Master of Science**, *Xi'an Jiaotong University*, Xi'an, Shaanxi, China.
Research Advisor: Prof. Qinghua Zheng, Department of Computer Science, Xi'an Jiaotong University
Major: Computer Science, GPA: 87.30/100.00
Master's Thesis: Constraint Solving Based Multithreaded Program Correctness Verification Method
- 2008 - 2012 **Bachelor of Science**, *Xi'an Jiaotong University*, Xi'an, Shaanxi, China.
Major: Software Engineering, GPA: 86.10/100.00

Work Experience

- Fall 2017 - Present **Graduate Research Assistant (Advisor: Vivek Sarkar)**, *Georgia Tech*, Atlanta, GA, USA.
Lechen's research focuses on the correctness and performance of accelerated applications, while exploring potential use cases of GPUs in speeding up dynamic analysis. Lechen works closely with researchers from national labs, LLVM community, and other universities.
- Fall 2020 **Compiler Intern**, *NVIDIA*, Hillsboro, OR, USA.
Implemented OpenMP Tool Interface (OMPT) for NVIDIA's OpenMP runtime. Reported a number of OMPT design issues to OpenMP's OMPT subcommunity.
- Summer 2020 **SDE Intern**, *Meta*, Menlo Park, CA, USA.
Participated in the development of a performance analysis tool for PyTorch applications. Added a number of new features into the tool and improved the tool's performance by 14.2 times. The optimized tool, including all new features, has been deployed in Meta.
- Summer 2019 **Visiting Scholar in IT Center (Collaborator/Supervisor: Joachim Protze, Christian Terboven, Matthias Müller)**, *RWTH Aachen University*, Aachen, Germany.
Developed Arbalest, a dynamic data mapping issue detector for OpenMP applications. Compared to other well-known concurrency bugs, data mapping issue is little-studied. Arbalest is the first dynamic analysis tool capable of detecting data mapping issues in OpenMP applications.

- Summer 2018 **Visiting Scholar in IT Center (Collaborator/Supervisor: Joachim Protze, Christian Terboven, Matthias Müller)**, *RWTH Aachen University*, Aachen, Germany.
Participated in two projects: a) design and implement DataRaceOnAccelerator (DRACC) micro-benchmark suite, and b) examine the implementation of OpenMP target offloading features in LLVM 8.0, reporting a number of bugs and inconsistencies to the LLVM community.
- Summer 2017 **Summer Research Fellow (Supervisor: Oscar Hernandez)**, *Oak Ridge National Laboratory (ORNL)*, Oak Ridge, TN, USA.
Collaborated with application developers at ORNL to find out potential concurrency bugs in OpenACC applications.
- Summer 2016 **Summer Research Fellow (Supervisor: Oscar Hernandez)**, *Oak Ridge National Laboratory (ORNL)*, Oak Ridge, TN, USA.
Participated in two projects: a) explore concurrency bug detection algorithms and tools for GPU applications, and b) analyze Julia's compiler to figure out the lowering process from Julia AST to LLVM IR.
- Fall 2015 – Spring 2017 **Graduate Research Assistant (Advisor: Vivek Sarkar)**, *Rice University*, Houston, TX, USA.
Research work was focused on concurrency bug detection and performance analysis for asynchronous many-task applications.
- Sep 2013 – June 2015 **Part-Time Software Development Engineer**, *Campanano Information Technology Co., Ltd.*, Xi'an, Shaanxi, China.
Lechen joined Campanano as a founding member and was the chief Java programmer in the team. Lechen participated in the development and maintenance of Easy-Lab digital lab management solution, including the database, the server-side code, and the client-side web pages.

Publications

- [1] Feiyang Jin, **Lechen Yu**, Tiago Cogumbreiro, Jun Shirako, Vivek Sarkar. Dynamic Determinacy Race Detection for Task-Parallel Programs with Promises. European Conference on Object-Oriented Programming (ECOOP) 2023

This paper elucidates the relationship between data races and program determinacy for task-parallel programs with *promises*, a common synchronization construct available in mainstream parallel programming frameworks.
- [2] **Lechen Yu**, Feiyang Jin, Joachim Protze, Vivek Sarkar. Leveraging the Dynamic Program Structure Tree to Detect Data Races in OpenMP Programs. IEEE/ACM International Workshop on Software Correctness for HPC Applications (Correctness) 2022

This paper proposes a new data race detector, TSAN-SPD3, for OpenMP programs. TSAN-SPD3 takes OpenMP tasks into account when detecting data races. Compared to Archer, the state-of-the-art OpenMP race detector, TSAN-SPD3 can mitigate false negatives without incurring a significant overhead.
- [3] **Lechen Yu**, Joachim Protze, Oscar Hernandez, Vivek Sarkar. ARBALEST - Dynamic Detection of Data Mapping Issues in OpenMP Applications. IEEE International Parallel and Distributed Processing Symposium (IPDPS) 2021: 464-474

In this paper we talk about data mapping issue, a concurrency bug incurred by incorrect usage of OpenMP's target offloading constructs. We present a formal definition of data mapping issue and introduce ARBALEST, a dynamic data mapping issue detector for OpenMP applications.

- [4] **Lechen Yu**, Joachim Protze, Oscar Hernandez, and Vivek Sarkar. A Study of Memory Anomalies in OpenMP Applications. International Workshop on OpenMP (IWOMP) 2020: 328-342

This paper summarizes our study on OpenMP's memory model. We present a group of memory anomalies OpenMP applications may encounter and analyze the root cause of these issues.

- [5] Adrian Schmitz, Joachim Protze, **Lechen Yu**, Simon Schwitanski, Matthias Mueller. DataRaceOnAccelerator - A Micro-benchmark Suite for Evaluating Correctness Tools Targeting Accelerators. HeteroPar 2019

In this paper we present DataRaceOnAccelerator (DRACC), a micro-benchmark suite covering common error patterns in accelerated applications. We introduce the design of DRACC and talk about included micro-benchmarks. We also illustrate the evaluation of four debugging tools using DRACC.

- [6] **Lechen Yu**, Vivek Sarkar. GT-Race: Graph Traversal Based Data Race Detection for Asynchronous Many-task Parallelism. Euro-Par 2018: 59-73

This paper introduces GT-Race, an efficient dynamic data race detector for asynchronous many-task applications.

- [7] Xiaodong Zhang, Zijiang Yang, Qinghua Zheng, Yu Hao, Pei Liu, **Lechen Yu**, Ming Fan, Ting Liu. Debugging Multithreaded Programs as if They Were Sequential. IEEE Access 6: 40024-40040

This paper is the extended version of [9] which introduces a formal behavior model for multithreaded programs. In this paper we elaborate each component of the proposed behavior model. Furthermore, we present the evaluation result of applying the behavior model to concurrency bug detection on a larger set of benchmarks.

- [8] Nick Vrvilo, **Lechen Yu**, Vivek Sarkar. A Marshalled Data Format for Pointers in Relocatable Data Blocks. ISMM 2017

This paper talks about a marshalled data format designed for pointers in relocatable data blocks. In addition, this paper also introduce how we implement the marshalled data format in Open Community Runtime, an open-source asynchronous many-task runtime.

- [9] Xiaodong Zhang, Zijiang Yang, Qinghua Zheng, Yu Hao, Pei Liu, **Lechen Yu**, Ming Fan, Ting Liu. Debugging Multithreaded Programs as if They Were Sequential. SATE 2016

In this paper, we present a formal behavior model for multithreaded programs. We also describe how to leverage this behavior model to detect specified concurrency bugs

- [10] Haijun Wang, Xiaohong Guan, Qinghua Zheng, Ting Liu, Xiangyang Li, **Lechen Yu**, Zijiang Yang. Reducing Test Cases with Causality Partitions. SEKE 2014

This paper proposes an optimal test case selection method that help programmers filter out unnecessary test cases in testing.

Presentations

March 2024 **SIAM Conference on Parallel Processing for Scientific Computing (PP24).**

Title: Determinacy Checking of Heterogeneous Parallel Programs

Feb 2023 **OpenMP Users Monthly Teleconferences.**

Title: Data Consistency Debugging for OpenMP Target Offload

- Feb 2020 **2020 Exascale Computing Project Annual Meeting.**
Title: Dynamic Data Mapping Issue Detection in OpenMP Applications
- June 2019 **OpenMP Tool Interface (OMPT) Regular Remote Meeting.**
Title: OpenMP Diagnostic Tool for Device Offloading
- Nov 2018 **The Fifth Workshop on the LLVM Compiler Infrastructure in HPC, Dallas, TX.**
Title: A Study of OpenMP Device Offloading in LLVM: Correctness and Consistency
Co-located with SC 18
- Nov 2018 **2018 CRNCH Summit, Atlanta, GA.**
Title: Extensions to OpenMP Device Offloading for Improved Performance and Debugging
- Mar 2018 **Summit Application Readiness Workshop 2018, Oak Ridge, TN.**
Title: Detecting and Fixing Data Race/Consistency Bugs
- Nov 2017 **Kokkos Hackathon at Georgia Tech, Atlanta, GA.**
Title: Lessons Learned from Porting BFS and Stencil-2D to Kokkos

Research Projects (selected)

- 2022 – **Improving the Performance of Data Race Detection for Structured Task Parallelism, Georgia Tech/RWTH Aachen University.**

Inspired by the task-oriented execution model in the OpenMP spec, this project extends SPD3, a data race detection algorithm designed for async-finish task parallelism to support OpenMP programs. By extending SPD3's key data structure, the Dynamic Program Structure Tree (DPST), SPD3 can support the majority of OpenMP constructs. We have implemented a prototype, TSan-spd3, on top of Google's ThreadSanitizer (TSan), and conducted an apples-to-apples comparison with Archer, the state-of-the-art dynamic race detector for OpenMP programs. The evaluation results on nine benchmarks from BOTS and SPEC OMP2012 show that in eight out of nine benchmarks TSan-spd3 achieved similar overhead with Archer, while TSan-spd3 can identify more potential races than Archer. [2]

Link <https://github.com/lechenyu/TSan-spd3>

Technique C/C++, OpenMP

Tool Chain LLVM, ThreadSanitizer, Archer

- 2019 – **Data Mapping Issue Detection in Accelerated OpenMP Applications, Georgia Tech/RWTH Aachen University.**

OpenMP is a popular parallel framework and widely applied in CPU applications. Starting from version 4.0, OpenMP added support for accelerators (e.g., Nvidia Tesla GPU, AMD Radeon GPU) by *target offloading* constructs. Since incorrect usage of target offloading constructs may introduce concurrency bugs, this project focuses on *data mapping issue*, a concurrency bug caused by erroneous data movement between the host and accelerators. We proposed a precise and sound dynamic detection algorithm for data mapping issues with $O(1)$ space and time overhead. We also designed Arbalest, a dynamic data mapping issue detector, Arbalest, as an extension to Archer, a state-of-art data race detector for OpenMP applications. [3]

Link <https://github.com/lechenyu/Arbalest>

Technique C/C++, OpenMP, CUDA

Tool Chain LLVM, ThreadSanitizer, Archer

- 2019 – 2021 **Accelerating Dynamic Analysis by GPUs**, *Georgia Tech/UMass Boston*.
 Dynamic analysis is wildly used in debugging and performance analysis, but it may incur huge time and space overhead to the program execution, which may further affect the program behavior. To reduce interference to the program execution, this project tries to propose an efficient dynamic analysis method that decouple the analysis procedure from the runtime data collection. The collected data is offloaded to other computing resources such as GPUs to carry out specified dynamic analysis (e.g., performance analysis, bug detection). With the massive parallelism on GPUs, the dynamic analysis can be significantly accelerated. Furthermore, the separation of data collection and analysis also reduces the overhead to the program execution.
- Technique Java, C/C++, CUDA
- Tool Chain Hotspot JVM, RoadRunner, LLVM, ThreadSanitizer, NVProf
- 2018 – 2019 **DataRaceOnAccelerator (DRACC) Benchmark Suites**, *Georgia Tech/RWTH Aachen University*.
 DRACC is a benchmark suite designed for evaluating the effectiveness of debugging tools in accelerated applications. DRACC consists of a group of benchmarks written in OpenMP, OpenACC, and CUDA, each of which represents a common error pattern related to memory access issues in accelerated applications. We collected these error patterns by comprehensively analyze accelerated applications from open-source benchmark suites (e.g., SPEC-ACCEL, Rodinia, DataRaceBench). With DRACC, we evaluated four debugging tools including ThreadSanitizer, Archer, CUDA-Racecheck, and GPUVerify to compare their precision and recall on collected error patterns. In addition, DRACC has been used by other work, such as OMPSan, to demonstrate the correctness of the proposed debugging tool. [5]
- Link <https://github.com/RWTH-HPC/DRACC>
- Technique C/C++, OpenMP, OpenACC, CUDA
- Tool Chain LLVM, ThreadSanitizer, Archer, CUDA-Racecheck, GPUVerify
- 2016 – 2018 **Data Race Detection for Asynchronous Many-task Parallelism**, *Rice University*.
 Designed an efficient data race detection algorithm for Open Community Runtime (OCR), an open-source asynchronous many-task (AMT) runtime developed by Intel, and implemented GT-Race, a prototype dynamic race detector for OCR applications. Considering the substantial number of tasks in an OCR application, GT-Race leverages the computational graph to store happens-before relations. Compared to classic vector-clock based race detectors, GT-Race significantly reduces memory overhead, while the time overhead is acceptable. In addition, the computational graph based race detection algorithm can also be applied to other AMT runtimes such as Cilk and Legion. [6]
- Technique C/C++, Intel Open Community Runtime
- Tool Chain Intel Pin
- 2016 – 2017 **Marshaled Data Format for Pointers in Relocatable Data Blocks**, *Rice University*.
 Designed a marshalled data format for relocatable data blocks to simplify the programming of data movement in a distributed application. The data format has been applied in ocxxr, a C++ library providing concise and C++-friendly wraps for existing APIs in Open Community Runtime (OCR) programming model. [8]
- Technique C/C++, Intel Open Community Runtime

2014 – 2015 **Constraint Construction and Optimized Solving For Intelligent Testing of Multithreaded Programs**, *Xi'an Jiaotong University*.

Designed a formal behavior model for multithreaded applications written in C/C++. The behavior model encodes the semantics of a multithreaded program, including happens-before relations and data-flows, into first order logic formulas. By means of efficient SMT solvers such as Z3, the behavior model has been utilized to detect multiple kinds of concurrency bugs such as data races and deadlocks. [7],[9]

Technique C/C++

Tool Chain LLVM, KLEE, Z3 SMT Solver

2013 – 2014 **Test Case Generation and Optimal Selection By Control-flow and Data-flow Analysis**, *Xi'an Jiaotong University*.

Designed a test case selection method to achieve the same path coverage as the original set of test cases while using minimum number of test cases. The test case selection method was implemented in Java Pathfinder, a popular verification tool for Java programs, as a bytecode analysis pass. The bytecode analysis pass calculates the control-flow and data-flow of a Java program, marks the causality between instructions, and finally guides Java PathFinder to avoid generating redundant test cases. [10]

Technique Java, JVM

Tool Chain Java Path Finder, Soot

Industry Projects

Sep 2013 – **Easy-Lab Digital Lab Management Solution**, *Campnano Information Technology Co., Ltd.*, Xi'an, Shaanxi, China.
June 2015

Easy-Lab is a commercial lab management platform which has been applied by multiple Chinese universities and research institutions. The key functionalities of Easy-Lab include facility and resource management, facility reservation, data analysis, report generation, and employee productivity statistics. As the key developer of Easy-Lab, Lechen tackled the database, the server-side code, and the client-side web pages.

Link <http://www.ylab.cn>

Technique Java, Hibernate, Spring, Struts, Javascript, JQuery, HTML, CSS

Tool Chain Tomcat, Nginx, SQL Server

Peer Review for Conferences

August 2018 **International Workshop on Languages and Compilers for Parallel Computing, LCPC 2018.**

External reviewer for conference proceedings

Professional Service and Activities

October 2021 **Mentor for ECP OpenMP Virtual Hackathon.**

Provided OpenMP and GPU expertise to a lattice QCD library. Helped the developers to implement atomic operations and reductions correctly using OpenMP constructs

May 2019 **Mentor for GPU Hackathon @ BNL.**

Provided CUDA and OpenMP expertise to an aerospace application (a pseudo-spectral direct numerical simulation application using large 3D fourier transforms) and helped engineers from IBM locate a bug in IBM XL compiler

Nov 2017 **Scribe of a BoF session (Memory-Centric Clouds and HPC) for International Conference for High Performance Computing, Networking, Storage and Analysis, SC 2017.**

Served as a scribe to record opinions and comments expressed by panelists and audiences, and helped panelists summarize the discussion after the BoF session.

May 2017 **Volunteer for ACM International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS 2017.**

Served as an interpreter to help the local arrangement committee communicate with attendees and led a group of volunteers during the conference excursion.

Teaching Experience

Spring 2023 **Lead Teaching Assistant**, *Georgia Tech*, Atlanta, GA, USA.

- Spring '23 CS 6241 (Graduate-Level Compiler Design); Tech: LLVM

Fall 2016 – **Teaching Assistant**, *Rice University*, Houston, TX, USA.

Spring 2017

- Fall '16 Comp 311 (Functional Programming); Tech: Scala

- Spring '17 Comp 322 (Fundamentals of Parallel Programming); Tech: Java

Travel Grant

July 2018 **International Symposium on Software Testing and Analysis, ISSTA 2018.**

Languages

English Fluent

Mandarin Fluent

Personal Information

Citizen of China. Visa Status: F1