

LARISA C. STOLTZFUS

Informatics Forum · 10 Crichton Street · Edinburgh, UK · EH8 9AB

larisa.stoltzfus@ed.ac.uk

<http://homepages.inf.ed.ac.uk/s1147290/>

Nationality: USA and UK

Objective

Seeking a computational scientist position at a research institution to develop and improve tools and frameworks for HPC applications, in particular large-scale physical simulations.

Education

- **PhD Informatics** - University of Edinburgh, UK
Sep 2016 - Dec 2020 (Expected)
My research focuses on bridging the gap between **HPC** applications and cutting edge **compiler** technology. My goal is to bring **performance portability** to large-scale physical simulations without losing productivity. Using **LIFT**, a functional IR for parallel languages, and **stencil DSLs** I aim to build a methodology for **3D wave models** to be programmed productively and compiled down to hardware-agnostic, optimized code. Results from my first year's work were published in a paper at a top tier compiler conference (CGO) which won Best Paper and subsequent work has been published in a renowned code generation journal (TACO).
- **MSc Pervasive Parallelism** - University of Edinburgh, UK
Sep 2015 - Sep 2016
Performance, Portability and Productivity for Room Acoustics Codes
Developed new parallel framework for swapping in data abstractions and optimizations for room acoustics simulations. Results published in a notable audio effects conference (DAFx).
- **MSc High Performance Computing, with Distinction** - University of Edinburgh, UK
Sep 2012 - Sep 2013
Improving Job Node Allocation on HECToR in Various Topological Manifestations
Developed and compared various job placement algorithms using historical data from a Cray XE6 run on a simulator.
- **BA Physics** - Grinnell College, USA
Sep 2001 - Dec 2005
GPA 3.4

Scholarships

- Sep 2015 - Sep 2019 **CDT Pervasive Parallelism EPSRC Grant EP/L01503X/1**
- Sep 2001 - Dec 2005 **Grinnell College Merit Scholarship**

Conference, Workshop and Invited Presentations

- **McGill University Systems Group Seminar** - Feb 2020 in Montreal, Canada
"LIFTing the Abstraction Layer of Room Acoustics Simulations"
- **HiPEAC Summer School** - July 2019 in Fiuggi, Italy
"Performance, Portability and Productivity for 3D Wave Models" (Poster)
- **LLNL Summer Intern Poster Session** - August 2018 in Livermore, USA
"Predictive Modelling of Optimal Data Placement on GPUs" (Poster)
- **WOLFHPC Workshop at the International Conference for High Performance Computing, Networking, Storage, and Analysis** - Nov 2017 in Denver, USA
"A Modular Approach to Performance, Portability and Productivity for 3D Wave Models"
- **University of Glasgow Systems Group Seminar** - Oct 2017 in Glasgow, UK
"3D Wave Modelling in LIFT"

- **International Conference on Digital Audio Effects** - Sep 2017 in Edinburgh, UK
“Performance Portability for Room Acoustics Simulations” (Poster)
- **Women in HPC Workshop at ISC** - June 2016 in Frankfurt, Germany
“Performance, Portability and Productivity for Room Acoustics Codes”
- **Critical Blue Company Presentation** - Apr 2016 in Edinburgh, UK
“Performance, Portability and Productivity for Room Acoustics Codes”
- **Ikon Science User Group Meeting** - Nov 2011 in Houston, USA
“Writing Plugins For the External Interface in RokDoc”
- **Forum Labo et Biotech** - June 2008 in Paris, France
“Novel Detection Technology Within Electrophoretic Separations”
- **Analytical Research Forum** - July 2007 at the University of Strathclyde (Poster)

Artifact Evaluation Committes

- IA³ Workshop (SC) - November 2018
- ACM/IEEE International Symposium on Code Generation and Optimization - February 2019

Publications

- **L. Stoltzfus**, B. Hagedorn, M. Steuwer, S. Gorlatch and C. Dubach. *“Tiling Optimizations for Stencil Computations Using Rewrite Rules in Lift”* (2019) ACM Transactions on Architecture and Code Optimization.
- **L. Stoltzfus**, M. Emani, P. Lin and C. Liao. *“Data Placement Optimization in GPU Memory Hierarchy using Predictive Modeling”* (2018) Workshop on Memory Centric High Performance Computing (SC).
- M. Bari, **L. Stoltzfus**, P. Lin, M. Emani, C. Liao and B. Chapman. *“Is Data Placement Optimization Still Relevant On Newer GPUs?”* (2018) Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems Workshop (SC).
- B. Hagedorn, **L. Stoltzfus**, S. Gorlatch, M. Steuwer and C. Dubach. *“High Performance Stencil Code Generation with Lift”* (2018) International Symposium on Code Generation and Optimization. **Best Paper Award.**
- **L. Stoltzfus**, C. Dubach, M. Steuwer, A. Gray and S. Bilbao. *“A Modular Approach to Performance, Portability and Productivity for 3D Wave Models”* (2017) International Workshop on Domain-Specific Languages and High-Level Frameworks for High Performance Computing (SC).
- **L. Stoltzfus**, A. Gray, C. Dubach and S. Bilbao. *“Performance Portability for Room Acoustics Simulations”* (2017) International Conference on Digital Audio Effects.

Teaching Experience

- Autumn 2015 - Autumn 2019 - **Lab Assistant for HPC MSc Courses**,
University of Edinburgh
Topics include: MPI, OpenMP, R, Parallel Languages, Numerical Algorithms, HPC Architectures, Software Development, Performance Programming, Scalability
- Spring 2018 - **Tutor for Software Testing UG Course**, *University of Edinburgh*
- Spring 2017 - **Demonstrator for Embedded Systems MSc Course**, *University of Edinburgh*
- Autumn 2016 - **Tutor for Computing Systems UG Course**, *University of Edinburgh*
- Autumn 2004 - Spring 2005 - **Lab Assistant and Tutor for Intro Physics UG Course**,
Grinnell College
- Autumn 2002 - Autumn 2004 - **Lab Assistant and Tutor for Intro Computing UG Course**,
Grinnell College

Professional Work Experience

- May - September 2018 - **Computational Scholar at Lawrence Livermore National Lab**
A research and development institution for science and technology applied to national security

PREDICTIVE MODELLING FOR GPU DATA PLACEMENT

- Designed and developed workflow for collecting, processing and analyzing hardware counter data from a suite of scientific benchmarks
- Developed tree-based machine learning model for predicting optimal data placement for various scientific GPU applications
- Implemented run-time prediction software for collecting hardware counter data during run-time using CUPTI to predict best data placement
- Work published in two papers at Supercomputing 2018 workshops

- October 2013 - June 2015 - **HPC Engineer at Lloyd's of London**
An umbrella insurance company providing a market for syndicate companies to share risk

CAPITAL CALCULATION KERNEL DEVELOPMENT:

- Developed quantitative software and integration tests for Monte Carlo based capital models using Test-Driven Development (TDD) in a bespoke functional DSL
- Produced regular mathematical analysis and implementation of model with actuarial team

OPTIMIZATION:

- Performed analysis of bottlenecks in the model leading to 900% faster run times
- Wrote data aggregation script in Powershell leading to 40% improved performance in syndicate model runs

RESTFUL SERVICE API AND GUI:

- Added functionality to job scheduling and data management web service and client in C#/.NET using TDD

GRID ECOSYSTEM DEVELOPMENT AND MANAGEMENT:

- Implemented checkpoint and restart functionality for model runs using Powershell scripts
- Lead workstream to improve stability of the system, liaising with operations group and external consultants

- February 2010 - September 2013 - **Software Developer at Ikon Science**
An oil and gas exploration, quantitative interpretation and software company

3D SEISMIC SECTION VIEWER DEVELOPMENT:

- Designed new user interface prototype using docking framework from third-party libraries
- Implemented new functionality for: updating vectorized 3D viewer, utilizing international locales, adding trendlines with parameters to data and navigating data with customized JTree classes
- Developed external plugin interface with JNA data types to liaise with C plugins

LICENSING AND BUILD MAINTENANCE:

- Served as Licensing Manager for two years, maintaining web-based licensing database, issuing licenses, upgrading licensing test suites and managing FlexLM issues
- Wrote and updated ANT scripts for Continuous Integration and installers

- November 2005 - February 2010 - **Research Analyst at deltaDOT**

An Imperial College spin-out focusing on the development and application of high energy physics algorithms to biomolecular sample detection

DNA SEQUENCING APPLICATION DEVELOPMENT:

- Researched, developed and implemented mobility shifts, artifact discrimination, smoothing, noise-elimination and peak finding algorithms
- Analyzed and optimized sequence identification functionality
- Upgraded application to output data in standard DNA-sequence file format

BIOMOLECULAR ANALYSIS SOFTWARE PACKAGE DEVELOPMENT:

- Researched and implemented data manipulation algorithms (peak finding, peak smoothing, mass and concentration calibration) for capillary electrophoresis instruments
- Added functionality for batch sample processing, project history, concentration calibration, supplementary graphical gel imaging and peak-by-peak exportation and analysis
- Designed and implemented tutorial and setup wizards

GENERAL SOFTWARE DEVELOPMENT:

- Created backup and search applications for processed sample data using MySQL and C++
- Redesigned data acquisition software for microfluidics chip and updated to use multi-threading

- Summer 2005 - **Research Intern (REU) at Arctic Region Supercomputing Center**

“Correcting the Top Boundary Conditions for the UAF Ionospheric Model”

- Updated vertical ion diffusion in real-time ionosphere model and redesigned accompanying MySQL database

- Summer 2004 - **Research Intern (REU) at University of California Davis**

“Implementing Aftershocks Into A Slider-Block Model”

- Designed and implemented algorithms for accommodating aftershocks in a Monte-Carlo model of earthquakes in southern California

- Summer 2003 - **Undergraduate Research Participant at Grinnell College**

“Estimating Gamma-Rays Using a Monte-Carlo Simulation”

- Analyzed effectiveness of a gamma-ray Monte-Carlo model

- Summer 2002 - **Undergraduate Research Participant at Grinnell College**

“Data Mining Student Web Usage Logs”

- Investigated data mining as an approach to understanding student use of online courses

Computing Skills

- **Programming Languages:** C, Java, C#, C++, Scala, R, Python, Scheme, Perl, FORTRAN
 - Object-Oriented experience in Scala, Java, C# and C++
 - Scripting experience in bash, Powershell, R, Python and Perl
 - Functional experience in Scala and Scheme
 - Experienced in Test-Driven Development and Agile SCRUM
- **Source Control:** Git, SVN, SCM, CVS
- **Parallel Frameworks:** OpenCL, MPI, OpenMP, CUDA
- **Web Languages:** HTML, CSS, PHP
- **Databases:** SQL