

# Craig Mustard

craig@gmail.com

1-604-710-2882

**Education** PhD, Electrical and Computer Engineering,  
September 2013 - November 2020  
University of British Columbia

B.Sc Computing Science, Sept 2007 - April 2011  
Simon Fraser University, Burnaby, BC

## Skills

**Overview** A generalist with deep expertise throughout the system stack: parallel programming, computer architecture, performance evaluation, compiler engineering, operating systems, programming languages, front and back-end web services. Experienced in software development for research and commercial applications.

**In Detail**

- C, C++, SQL, Python, Java, Scala, Javascript/React, and more. Docker, AWS, Google Cloud, Azure, and Kubernetes. Git, code reviews, testing.
- In depth knowledge in areas such as: data analytics system internals, distributed systems, databases, compilers, concurrency, data analysis. Basic machine learning experience.
- Software design for high performance on novel computer architecture, low-level analysis on simulation and real hardware. Contributions to architecture designs and simulators.
- Experience leading systems research projects and software development. Technical writing, Project management. Enjoys working with others to build efficient, easy to use, and reliable systems that meet goals.

## Research and Industry Experience

Jan 2021 - Now Senior Member of Technical Staff, Oracle (Previously Oracle Labs). Working for HeatWave (MySQL Analytics) group developing HeatWave's cloud control plane.

Sept - Oct 2020 Consultant, Infrastructure at BC NDP. Developed and operated infrastructure and critical services to support election operations for central campaign and 87 riding-level campaigns during a challenging election period. (Kubernetes, RDS/Postgresql, SSO).

2013 - 2020 PhD supervised by Dr. Alexandra Fedorova (Candidate since 2017. Worked for 2 years during PhD.)

- Thesis: *Compilation-assisted Performance Acceleration for Data Analytics*. Research themes: novel computer architectures, data analytics, and compilers.
- Projects included: (1) Compiling and orchestrating data analytics tasks (SQL queries) to run on emerging in-network processors. Integrated with Apache Spark and Presto. (Published at HotCloud 2019) (2) Enabling transparent data sharing between users of data analytics systems by detecting equivalent code. Integrated with Apache Spark (Awarded best paper at IEEE BigData 2018). (3) Performing qualitative study on how developers understand performance, including open coding and analysis. (FSE 2018).

2015 - 2016 Senior Member of Technical Staff, Oracle Labs, Vancouver on RAPID project (Part-Time, 1 year 2 months)

- Developed workloads for novel computer architecture designed for data analytics. Built streaming operators for aggregate queries and a highly optimized JSON parser.
  - Workload-driven optimization on GCC compiler back-end to improve performance and enable better use of architecture specific instructions (RTL / C).
- 2014 - 2015      Research Assistant at Oracle Labs, Vancouver on RAPID project. (Full time, 1 year)
- Analyzed low level performance of database and big data workloads on a novel computer architecture. Developed performance improvements and presentations to inform team on the impacts of architectural design choices.
  - GCC compiler maintenance and improvements. Migrated team from 4.4.5 to 4.9.0, assisted with modifications to the compiler backends, fixed software issues, enabled C++ compiler. Enabled and evaluated performance of GCC's link time and profile guided optimizations on database software.
  - Developed several new features and reports for CI-driven performance analysis tool.
- Summer 2013      Research Internship at STMicroelectronics, SoC Platform Automation Group. (4 months)
- Performed research relating to the STHORM multicore processor with explicitly managed memory hierarchy (scratchpad+DMA). Evaluated improvements to ease programmability and performance of using automatic software cache and a custom data management library. Developed computer vision algorithm tests in OpenCL.
- 2011 - 2013      MSc Student supervised by Dr. Alexandra Fedorova. (Transferred to PhD in 2013)
- Researched how to improve programmability for processors with explicitly managed memory hierarchies. Developed automatic software caches, and access pattern aware libraries aware libraries for workloads such as graph analytics and MapReduce. (Work presented at OSDI2012, DATE2013, HotPar 2013).
- 2009 - 2011      Undergraduate Research Assistant under Dr. Alexandra Fedorova. (During B.Sc)
- Worked on Synchronization via Scheduling, which scheduled parallel accesses to complex data structures using bloom filters. Published in PLDI 2011.
  - Duties included: Developed a parallel task manager. Language design and implementation. Experiment development and comparison with Software Transactional Memory (Intel STM and Dresden STM).
- 2004 - 2007      Co-Founder of Clear-Site Web Solutions. Developed custom web applications, a re-brandable content management system, embedded Linux digital advertising signs, and operations infrastructure.

## Graduate Coursework

- UBC                      EECE 571B - Big Data Systems, EECE 527 - Advanced Computer Architecture
- SFU                      CMPT 886 - Special Topics in Multicore Architecture, CMPT 885 - Special Topics in Computer Architecture, CMPT 741 - Data Mining, CMPT 705 - Design and Analysis of Algorithms, CMPT 726 - Machine Learning.

## Teaching Experience

- Fall 2012              Teaching Assistant for Compiler course, using LLVM, flex and bison. (CMPT 379 taught by Anoop Sarkar, Simon Fraser University)

## Awards

- 2013, 2014 Graduate Fellowship (PhD), Simon Fraser University  
2009, 2010, 2011 Holder of three NSERC Undergraduate Research Awards

## Refereed Publications

Swati Goswami, Nodir Kodirov, **Craig Mustard**, Ivan Beschastnikh, Margo Seltzer *Parking Packet Payload with P4*, CoNEXT 2020

Joel Nider, **Craig Mustard**, Andrada Zoltan, Alexandra Fedorova *Processing in Storage Class Memory*, HotStorage 2020

**Craig Mustard**, Fabian Ruffy, Anny Gakhokidze, Alexandra Fedorova, Ivan Beschastnikh *Jumpgate: In-Network Processing as a Service for Data Analytics*, HotCloud 2019

**Craig Mustard**, Alexandra Fedorova *Practical Cross Program Memoization with Key-Chain*, IEEE BigData 2018.

Alexandra Fedorova, **Craig Mustard**, Ivan Beschastnikh, Julia Rubin, Augustine Wong, Svetozar Miucin, Louis Ye, *Performance comprehension at WiredTiger*, FSE 2018.

**Craig Mustard**, Alexandra Fedorova, Ivan Bestchastnickh *Jumpgate: Towards In-Network Data Processing*, OSDI 2018 (Poster).

**Craig Mustard**, Svetozar Miucin, Yuan Liu, Karol Swietlicki, Yuxuan Li, Alexandra Fedorova, Arrvindh Shriraman, *Access Declarations: Revealing Data Access Patterns to Hardware*, HotPar 2013 (Poster).

Mark Roth, Micah J Best, **Craig Mustard** and Alexandra Fedorova, *Deconstructing the Overhead in Parallel Applications*, IEEE International Symposium on Workload Characterization, 2012.

Micah J Best, Shane Mottishaw, Mark Roth, **Craig Mustard**, Alexandra Fedorova, Andrew Brownsword, *Synchronization via Scheduling: Techniques For Efficiently Managing Shared State in Video Games*, in 32nd ACM SIGPLAN conference on Programming Language Design and Implementation (PLDI 11).

Micah J Best, Shane Mottishaw, **Craig Mustard**, Mark Roth, Parsiad Azimzadeh, Alexandra Fedorova, Andrew Brownsword, *Schedule Data, Not Code*, in Third USENIX Workshop on Hot Topics on Parallelism (HotPar 11).

Micah J Best, Shane Mottishaw, **Craig Mustard**, Mark Roth, Alexandra Fedorova and Andrew Brownsword, *Synchronization via Scheduling: Managing Shared State in Video Games*, in Second USENIX Workshop on Hot Topics on Parallelism (HotPar 10).

Jon Hourd, Chaofei Fan, Jiasi Zeng, Qiang (Scott) Zhang, Micah J Best, Alexandra Fedorova and **Craig Mustard**, *Exploring Practical Benefits of Asymmetric Multicore Processors*, In Proceedings of the 2009 Workshop on Parallel Execution of Sequential Programs on Multi-core Architectures, in conjunction with ISCA-36.

Micah J Best, Alexandra Fedorova, Ryan Dickie, Andrea Tagliasacchi, Alex Couture-Beil, **Craig Mustard**, Shane Mottishaw, Aron Brown, Zhi Feng Huang, Xiaoyuan Xu, Nasser Ghazali and Andrew Brownsword, *Searching for Concurrent Design Patterns in Video Games: Practical lessons in achieving parallelism in a video game engine*, 15th International European Conference on Parallel and Distributed Computing (Euro-Par 2009), August 2009.

## Non-refereed work

Craig Mustard, Alexandra Fedorova, *Access declarations for explicitly managed memory*, Invited Talk at STHORM Workshop, DATE Conference 2013.

Craig Mustard, Alexandra Fedorova, Arrvindh Shriraman, *Evaluating Software Managed Memory with MapReduce* (Poster), in 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 12).

April 5, 2021