

# Saman Ashkiani



+1 (530)750-9038  
sa.ashkiani@gmail.com  
<https://sashkiani.github.io/>  
<https://github.com/sashkiani>  
<https://www.linkedin.com/in/saman-ashkiani-88656352/>

## INTEREST

My primary expertise is in general-purpose GPU computing (GPGPU). I got my PhD with an emphasis on primitive *parallel algorithms* and *dynamic concurrent data structures*, specifically designed and implemented for modern GPUs. Moreover, I get excited in dealing with performance related problems such as:

1. High performance computing (HPC), parallel and distributed systems
2. Compiler methods, code generation, JIT compilation (e.g., LLVM)
3. Database management systems, analytical processing
4. Machine learning, training/inference performance

## WORK EXPERIENCE

CURRENT, FROM MAY 2018 (FT)

OmniSci, Inc.  
**Backend Engineer**

Working on the backend query engine of OmniSciDB, a high performance GPU-accelerated relational database management system (RDBMS) primarily designed on the idea of real-time JIT compilation of SQL queries into GPU-friendly code. This work involves using C++, LLVM, and CUDA.

JAN 2018 – MAY 2018 (FT)

University of California, Davis  
**Postdoctoral Scholar**

Worked on the design and analysis of dynamic data structures for sparse data representation on the GPU, and more specifically in the context of dynamic graph storage and an eventual efficient processing of that data structure. This work involved C++ and CUDA developments.

JUNE 2016 – SEP 2016 (FT)

Adobe Research, Adobe Systems, Inc.  
**Data Scientist Intern**

Exploring the possibility of GPU acceleration for certain analytics operations on the Apache Spark framework. This work involved using Scala and Spark.

## EDUCATION

2011 – 2017 **Doctor of Philosophy**  
Electrical and Computer Engineering  
*University of California, Davis*

2015 **Master of Science**  
Electrical and Computer Engineering  
*University of California, Davis*

2006 – 2011 **Bachelor of Science**  
Electrical Engineering  
*Sharif University of Technology*

## AWARDS

2016 **NVIDIA Graduate Fellowship**  
2017 **Dissertation Writing Fellowship**  
*University of California, Davis*

## COMPUTER SKILLS

PROGRAMMING LANGUAGES	C/C++, Python, CUDA Java, Scala Haskell
COMPILER METHODS	LLVM
PARALLEL FRAMEWORKS	OpenMP, MPI
COMPUTER GRAPHICS	OpenGL, Vulkan
DBMS & ANALYTICS	SQL, MATLAB, R, Spark
MACHINE LEARNING	TensorFlow, Keras

## OPEN-SOURCE INVOLVEMENTS

OmniSciDB  
<https://github.com/omnisci/omniscidb>

SlabHash, a dynamic hash table for the GPU  
<https://github.com/owensgroup/SlabHash>

GPU Multisplit  
<https://github.com/owensgroup/GpuMultisplit>

## SELECTED PUBLICATIONS

**Saman Ashkiani**, Martin Farach-Colton, John D. Owens, “A Dynamic Hash Table for the GPU”, Proceedings of the 32nd IEEE International Parallel and Distributed Processing Symposium (IPDPS 2018).

**Saman Ashkiani**, Shengren Li, Martin Farach-Colton, Nina Amenta, John D. Owens, “GPU LSM: A Dynamic Dictionary Data Structure for the GPU”, Proceedings of the 32nd IEEE International Parallel and Distributed Processing Symposium (IPDPS 2018).

**Saman Ashkiani**, Andrew Davidson, Ulrich Meyer, John D. Owens, “GPU Multisplit: an extended study of a parallel algorithm”, ACM Transactions on Parallel Computing (TOPC 2017).

**Saman Ashkiani**, Andrew Davidson, Ulrich Meyer, John D. Owens, “GPU Multisplit”, Proceedings of the 21st ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP 2016).

**Saman Ashkiani**, Nina Amenta, John D. Owens, “Parallel Approaches to the String Matching Problem on the GPU”, In Proceedings of the 28th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA 2016).

Muhammad Awad, **Saman Ashkiani**, Rob Johnson, Martin Farach-Colton, John Owens, “Engineering a High-Performance GPU B-Tree”, Proceedings of the 24th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP 2019).