Suhail Basalama

Ph.D. Candidate, University of California, Los Angeles
Google Scholar | GitHub | LinkedIn | Website | basalamasuhail@gmail.com

RESEARCH INTERESTS

- AI/ML software-hardware co-design
- Spatial, systolic, and dataflow architectures
- Performance and system modeling
- Domain-specific and customized computing
- Compiler construction for AI/ML accelerators
- Design-space exploration methodologies
- Hardware compilation and synthesis
- EDA and VLSI design methodologies

EDUCATION

University of California, Los Angeles — September 2020 | June 2026 (expected)
Ph.D. in Computer Science, Cumulative GPA: 3.9 /4.0 — Advisor: Prof. Jason Cong

University of California, Los Angeles — September 2020 | December 2024
M.S. in Computer Science, Cumulative GPA: 3.9 /4.0 — Advisor: Prof. Jason Cong

University of Arkansas, Fayetteville — August 2017 | December 2019
B.S. in Computer Engineering, Cumulative GPA: 4.0 /4.0

University of Arkansas, Fayetteville — August 2017 | May 2020
B.A. in Political Science, Cumulative GPA: 4.0 /4.0

Hinds Community College, Raymond, Mississippi — January 2016 | May 2017
Cumulative GPA: 4.0 /4.0

ACADEMIC EXPERIENCE

University of California, Los Angeles, *VAST Lab* — June 2020 | Present
Graduate Student Researcher — Advisor: Prof. Jason Cong

Stream-HLS: An automatic end-to-end MLIR-based compiler that converts sequential multi-kernel applications into optimized parallel dataflow architectures with streaming capabilities

- Developed multiple MLIR passes for the global optimization of multi-kernel applications
- Constructed a comprehensive design space beyond those of prior works
- Devised a global scheduling approach using mixed-integer nonlinear programming (MINLP)
- Evaluated Stream-HLS against SoTA frameworks achieving considerable speedups
- Published as a conference paper at ***FPGA 2025*** **(Best Paper Candidate)**

Odyssey: A comprehensive design space exploration framework for systolic arrays using a combination of mathematical programming, an evolutionary algorithm, and a padding-based heuristic

- Analyzed the impact of choosing divisor vs. non-divisor tiling factors in systolic array designs
- Designed a near-optimal heuristic algorithm for pruning the design space and reducing search time
- Published as a conference paper at ***DAC 2023***

FlexCNN: An HLS framework for composing CNN accelerators on FPGA with versatile systolic arrays

- Designed versatile systolic arrays to support transposed and dilated convolution layers
- Improved the front-end of the framework to take ONNX CNN models as inputs
- Implemented VGG16, U-Net, and E-Net with comparable or better performance than prior works
- Published as a journal paper at ***TRETS 2023***

Versatile Systolic Arrays (VSA): An HLS versatile systolic array for efficient processing of transposed and dilated convolution on FPGAs

- Implemented a decomposition of transposed and dilated convolutions as a versatile systolic array
- Automated the process of generating new VSA designs based on the target CNN
- Published as a poster at ***FCCM 2022***

University of Arkansas, Computer Systems Design Laboratory — January 2019 | May 2020
Undergraduate Student Researcher — Advisor: Prof. David Andrews

SPAR-2 Project: FPGA-based SIMD serial systolic array processors for accelerating machine learning workloads

- Built a MicroBlaze System-on-Chip (SoC) with separate BRAM for instructions
- Wrote an Instruction Sequencer Module in Verilog and packaged it as an AXI4 IP core
- Manipulated Booth's and Modified Booth's algorithms to reduce the shift operations
- Published as conference papers at ***ICDIS 2020*** and ***FPL 2021***

University of Arkansas, Smart Embedded Systems Lab — June 2018 | January 2019
Undergraduate Student Researcher — Advisor: Prof. Christophe Bobda

ARLO Robot: An autonomous ground vehicle built using the Parallax Arlo Robot System and the Digilent Zybo Z7-20 FPGA

- Built a Zynq-7000 SoC with UART IP and a custom Ultrasonic Sensor IP I implemented
- Deployed a Petalinux system on the SoC and configured the kernel with the needed modules

Cryptocurrency Wallet: A hardware-based wallet with end-to-end AES encryption for cryptocurrency on the Lattice iCE40 Ultra Wearable Development Platform

- Helped develop a UART-Bluetooth communication interface between the FPGAs and a smartphone
- Helped implement the AES encryption algorithm in Verilog on FPGA

INDUSTRY EXPERIENCE

Micron, Catapult-MLIR — June 2023 | September 2023
Intern - AI/ML Algorithms, Silicon Systems AI Managers: Dr. Nihar Athreyas and Dr. Lorenzo Ferretti

Developed an MLIR infrastructure for automatic code transformation for ASIC designs using Catapult-Siemens high-level synthesis

Microsoft, Brainwave Project — June 2021 | September 2021
Silicon/Firmware Intern Researcher — Managers: Dr. Steven K. Reinhardt and Dr. Gabriel Weisz

Built a command-line interactive performance visualization/profiling tool that analyzes RTL waveforms and constructs a visualization of the task dependencies and their run-time traces

PUBLICATIONS

[**ASP-DAC'26**] Stefan Abi-Karam, Rishov Sarkar, **Suhail Basalama**, Jason Cong, Callie Hao. "FIFOAdvisor: A DSE Framework for Automated FIFO Sizing of High-Level Synthesis Designs." To appear at *2026 ACM Proceedings of the 31st Asia and South Pacific Design Automation Conference (ASP-DAC)*

[**FCCM'25**] Huifeng Ke, Sihao Liu, Licheng Guo, Zifan He, Linghao Song, **Suhail Basalama**, Yuze Chi, Tony Nowatzki, and Jason Cong. "NoH: NoC Compilation in High-Level Synthesis." In *2025 IEEE 33rd Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM)*, pp. 162-171. IEEE, 2025. Open Source: https://github.com/rapidstream-org/rapidstream-noc

[**FPGA'25**] **Suhail Basalama**, and Jason Cong. "Stream-HLS: Towards Automatic Dataflow Acceleration." In *Proceedings of the 2025 ACM/SIGDA International Symposium on Field Programmable Gate Arrays (FPGA)*, 2025. **(FPGA'25 Best Paper Candidate)**. Open Source: https://github.com/UCLA-VAST/Stream-HLS

[**ASPLOS'24**] Neha Prakriya, Yuze Chi, **Suhail Basalama**, Linghao Song, and Jason Cong. "TAPA-CS: Enabling Scalable Accelerator Design on Distributed HBM-FPGAs." In *Proceedings of the 29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS), Volume 3*, pp. 966-980. 2024. Open Source: https://github.com/UCLA-VAST/TAPA-CS

[**FPGA'23**] Linghao Song, Licheng Guo, **Suhail Basalama**, Yuze Chi, Robert F. Lucas, and Jason Cong. "Callipepla: Stream Centric Instruction Set and Mixed Precision for Accelerating Conjugate Gradient Solver." In *Proceedings of the 2023 ACM/SIGDA International Symposium on Field Programmable Gate Arrays (FPGA)*, pp. 247-258. 2023. Open Source: https://github.com/UCLA-VAST/Callipepla

[**DAC'23**] **Suhail Basalama**, Jie Wang, and Jason Cong. "A Comprehensive Automated Exploration Framework for Systolic Array Designs." In *2023 60th ACM/IEEE Design Automation Conference (DAC)*, pp. 1-6. IEEE, 2023. Open Source: https://github.com/UCLA-VAST/Odyssey

[**TRETS'23**] **Suhail Basalama**, Atefeh Sohrabizadeh, Jie Wang, Licheng Guo, and Jason Cong. "FlexCNN: An end-to-end Framework for Composing CNN Accelerators on FPGA." *ACM Transactions on Reconfigurable Technology and Systems (TRETS) 16*, no. 2 (2023): 1-32. https://github.com/UCLA-VAST/FlexCNN

[**FCCM'22**] **Suhail Basalama**, Atefeh Sohrabizadeh, Jie Wang, and Jason Cong. "A Versatile Systolic Array for Transposed and Dilated Convolution on FPGA." In *2022 IEEE 30th Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM)*, pp. 1-2. IEEE, 2022.

[**FPL'21**] Atiyehsadat Panahi, **Suhail Basalama**, Ange-Thierry Ishimwe, Joel Mandebi Mbongue, and David Andrews. "A Customizable Domain-specific Memory-centric FPGA Overlay for Machine Learning Applications." In *2021 31st International Conference on Field-Programmable Logic and Applications (FPL)*, pp. 24-27. IEEE, 2021. Open Source: https://github.com/SuhailB/ArrayProcessor

[**ICDIS'20**] **Suhail Basalama**, Atiyehsadat Panahi, Ange-Thierry Ishimwe, and David Andrews. "SPAR-2: A SIMD Processor Array for Machine Learning in IoT Devices." In *2020 3rd International Conference on Data Intelligence and Security (ICDIS)*, pp. 141-147. IEEE, 2020. Open Source: https://github.com/SuhailB/ArrayProcessor

TEACHING EXPERIENCE

University of California, Los Angeles, Digital Design CS M152A — March 2023 | June 2023
Teaching Assistant — Lecturer: Prof. Majid Sarrafzadeh

University of California, Los Angeles, Digital Design CS M152A — March 2022 | June 2022
Teaching Assistant — Lecturer: Prof. Majid Sarrafzadeh

University of Arkansas, Digital Design CSCE 2114 — August 2019 | December 2019
Teaching Assistant — Lecturer: Prof. Patrick Parkerson

AWARDS

2025 Best Paper Candidate at International Symposium on Field-Programmable Gate Arrays 2025
2019 The Foundation for the International Exchange of Students Scholarship at UARK
2019 Dr. Henry M. Alexander Memorial Award
2019 Rosecrans, Sr Endowed Memorial Scholarship
2018 The Charles D. Brock Scholarship by the College of Engineering at UARK
2018 The Foundation for the International Exchange of Students Scholarship at UARK
2017 The John and Marie Lavallard International Scholarship at the University of Arkansas
2017 The University of Arkansas Transfer Student Scholarship
2014 The Silver Medal Representing Yemen in the Third Gulf Mathematics Olympiads in Oman
2013 The Top-Ten Student Ministerial Scholarship from the Yemeni government
2013 Ranked 9th among 218,964 Yemeni students in the National High-school Exams ***(Top 0.000041%)***

TALKS

- September 2025 Guest Lecture at **Yale University**: "Dataflow Hardware Acceleration of DNNs with High-Level Synthesis." **YouTube Link**
- July 2024 Member Talk at **SpatialML**: "End-to-End CNN Acceleration and Systolic Array Design Space Exploration for FPGAs." **YouTube Link**

TECHNICAL SKILLS

- **Design Skills**: Performance Modeling and Profiling, Design Space Exploration, System-on-Chip, IP Packaging, Static Timing Analysis, Algorithms, Finite State Machines
- **Languages**: C, C++, Python, High-level synthesis (HLS), Verilog, VHDL, Java, Tcl, Bash
- **CAD Tools**: Vitis HLS, Vivado Xilinx/AMD, Catapult HLS, Modelsim, Synopsys Design Compiler
- **Technologies**: FPGAs (Xilinx/AMD, Altera Intel, Lattice), Raspberry Pi, Arduino
- **Operating** Systems: Linux, Windows, macOS
- **Other** Skills: Git, SVN, LaTeX, Microsoft Office Suite

REFERENCES

Prof. Jason Cong, Distinguished Professor
Director, Center for Customizable Domain-Specific Computing
Director, VLSI Architecture, Synthesis, and Technology (VAST) Laboratory
CS and ECE Departments
University of California, Los Angeles, Los Angeles, California, USA
Phone: (310) 206-2775
E-mail: cong@cs.ucla.edu
Scholar Profiles: Personal Page | Google Scholar

Prof. David Andrews, Professor
Director, Computer Systems Design Laboratory
EECS Department
University of Arkansas, Fayetteville, Arkansas, USA
Phone: (479) 575-6529
E-mail: dandrews@uark.edu
Scholar Profiles: Personal Page | Google Scholar

Prof. Christophe Bobda, Professor
Director, SmartSystems Lab
Citi Endowed Professor in Advanced Technologies
ECE Department
University of Florida, Gainesville, Florida, USA
Phone: (352) 294-2024
E-mail: cbobda@ece.ufl.edu
Scholar Profiles: Personal Page | Google Scholar