Houman Homayoun

Email: hhomayou@gmu.edu
Web: http://ece.gmu.edu/~hhomayou/
https://nsfchest.org/
Lab: ASEEC: Accelerated, Secure, and Energy-Efficient Computing Lab
Department of Electrical and Computer Engineering
Department of Computer Science (Courtesy Appointment)
Department of Information Sciences and Technology (Courtesy Appointment)
George Mason University
Office: 3223, Engineering Building, Phone: (703) 993-5430

EDUCATION

- **Postdoc** *Sept. 2010-Aug. 2012*

Department of Computer Science and Engineering, University of California, San Diego
  Mentor: Prof. Dean Tullsen

- **PhD** *Sept. 2006-Sept. 2010*

Department of Computer Science, University of California, Irvine.
  ➢ Thesis: Beyond Memory Cells for Leakage and Temperature Control in SRAM-based Units, the Peripheral Circuits Story.
  Advisors: Prof. Alex Veidenbaum, Prof. Jean-Luc Gaudiot, Prof. Fadi Kurdahi

- **Master of Applied Science** *September 2003-March 2005*

Electrical and Computer Engineering Department, University of Victoria, Canada.
  ➢ Thesis: Using Lazy Instruction Prediction to Reduce Processor Wakeup Power Dissipation.

- **Bachelor of Science** *October 1998-May 2003*

Electrical and Computer Engineering Department, Sharif University of Technology.

EMPLOYMENT

- *Associate Professor*, George Mason University, Department of Electrical and Computer Engineering, Courtesy Appointment with the Department of Computer Science, Courtesy Appointment with the Department of Information Sciences and Technology *August 2018-present*

- *Advisory Committee, Research and Technology Commercialization (R&TC), Cybersecurity working group, Commonwealth of Virginia* *May 2018 - present*

- *Assistant Professor*, George Mason University, Department of Electrical and Computer Engineering, Courtesy Appointment with the Department of Computer Science, Courtesy Appointment with the Department of Information Sciences and Technology. *Aug. 2012-August 2018*

- *Board of Advisory Member*, BroadPak Corporation, Santa Clara, California, USA. *July 2012-Present*

- *NSF/CCC-CRA Computing Innovation Fellow*, University of California San Diego, Department of Computer Science and Engineering (Mentor: Dean M. Tullsen) *Sept. 2010-Aug. 2012*

- *Graduate Research Assistant*, University of California, Irvine, Department of Computer Science (Advisors: Alex Veidenbaum, Jean-Luc Gaudiot and Fadi Kurdahi) *Sept. 2006-Sept. 2010*

- *Design Architect*, Novelics Inc., Aliso Viejo, California, USA. *Jan. 2007-Oct. 2008*

- *Researcher Assistant*, McMaster University, Canada, Department of Electrical and Computer Engineering. *Oct. 2005-Apr. 2006*

- *Graduate Research Assistant*, University of Victoria, Canada, Department of Electrical and Computer Engineering. *Sept. 2003-Mar. 2005*

- *Research Assistant*, Sharif University Technology, Tehran, Electronic Research Center. *Oct. 2002-Jun. 2003*

GRANTS

Sponsored Research: Total $7,295,000, GMU Portion: $5,339,000, My Share: $4,395,000.

- *"Planning IUCRC George Mason University: Center for Hardware and Embedded System Security and Trust (CHEST)"* https://nsfchest.org/ *2018-2023*

NSF IUCRC, **(PI), $15,000 for 2017-2018 planning and $750,000 (expected) over 5 years for center expenses during phase I.**
Role: PI and Center Director on GMU site, A collaborative effort composed of George Mason University, Northeastern University, University of Connecticut, University of Texas at Dallas, University of Virginia, and Wright State University to establish the first NSF/AFRL center on HW and Embedded Systems Security and Trust.

- *"Obfuscated Logics to Enhance Security and Prevent Reverse Engineering"*   *2017-2020*
  DARPA MTO Office, **(PI), $1,800,000.** ($600K fab cost to GF), **PM: Kerry Bernstein**
  Role: Team lead on design and fabricating obfuscated logics in 14nm with GlobalFoundries.
- *"Mobilizing the Micro-Ops: Securing Processor Architectures via Context Sensitive Decoding"*   *2018-2021*
  DARPA MTO Office, SSITH program **(PI on GMU site), Total: $1,200,000. GMU share ($400,000), PM: Linton Salmon**
  Role: Leading the team to detect HW vulnerabilities in out-of-order processors
- *"Evolution of Computer Vision for Low Power Devices, Breaking its Power Wall and*   *2017-2020*
  Computational Complexity"
  NSF CSR-CNS, **(Co-PI)**, **$500,000**.
  Role: Developing an approximate Iterative Convolutional Neural Network coprocessor that supports approximation in memory and logic.
- *"3D-Split of Obfuscation and Authentication of logic"*   *2018-2019*
  DARPA MTO Office, **(Co-PI)**, **$495,000. PM: Ken Plaks**
  Role: Developing 3D-SOUL secure-compiler for cell, route and FSM obfuscation.
- *"Persistence and Extraction of Digital Artifacts from Embedded Systems"*   *2016-2017*
  NIST, National Cybersecurity Center of Excellence,
  (Co-PI), $75,000.
  Role: Establishing the persistence of digital artifacts on embedded systems through JTAG analysis.
- *"Hybrid Spin Transfer Torque-CMOS Technology to Prevent Design Reverse Engineering"*   *2015- 2017*
  DARPA MTO Office, **(PI), $349,000. PM: Kerry Bernstein**
  Role: Directing the project to design and deploy new circuit methods to build reconfigurable logics to enhance performance and power efficiency.
- *"Heterogeneous Ultra Low Power Accelerator for Wearable Biomedical Computing"*   *2015- 2018*
  NSF CSR-CNS, **(PI),** Total **$500,000,** GMU portion **$288,000.**
  Role: Directing the project to design and deploy new circuit methods to build reconfigurable logics to enhance performance and power efficiency.
- *"A Novel Biomechatronic Interface Based on Wearable Dynamic Imaging Sensors"*   *2013- 2018*
  NSF CPS – CNS, **(Co-PI) $995,000.**
  Role: Designing a heterogeneous architecture for computing intensive biomedical application, Compare with state-of-the-art heterogeneous platforms such as TI OMAP and Nvidia Tegra.
- *"Enhancing the Security on Embedded Automotive Systems"*   *2013- 2016*
  General Motors, **(Co-PI) $261,000.**
  Role: Hacking the CAN Bus Network of GM Cars.
- *"Inter-core Selective Resource Pooling in a 3D Chip Multiprocessor"*   *2010- 2012*
  NSF CI Fellow Award, NSF 1019343/CRA Sub Award CIF-B-68, **(PI), $280,000**.

Equipment Support from Industry
- Nvidia Corporation: 2 Tesla K40 GPU for CNN training, **$9,560**   *2017*
- Xilinx Corporation, 12 Xilinx ZYNQ board for HW accelerated computer vision, **$5,940**   *2016*
- Intel Corporation, 20 Intel Galileo and Intel Edison board for wearable computing, **$2,170**   *2015*

RESEARCH INTEREST
- **Computer System Cybersecurity (Current)**
  - o Online malware detection
  - o Adversarial machine learning
  - o Side-channel processor architecture defense and attack
  - o Detecting and containing malware epidemic in IoT network

- o   Reverse engineering
- **Big Data Computing (Current)**
  - o   Algorithms for energy-efficient acceleration of Big Data
  - o   Deep machine learning and data mining acceleration on heterogeneous platforms
  - o   Applied machine learning for cloud workload management, scheduling and tuning
  - o   Emerging big data application benchmarking and characterization on heterogeneous architectures
  - o   Mapping data and model parallel big data frameworks to heterogeneous accelerator architectures
- **Heterogeneous Architecture Design and Management (Current)**
  - o   Design space exploration of FPGA+CPU architecture for emerging big data frameworks
  - o   Scheduling and resource management in heterogeneous multicore CPU+FPGA architectures
  - o   Accelerator design for wearable biomedical applications
  - o   3D dynamic heterogeneous architecture design
- **Emerging Memory Technologies (Current)**
  - o   Emerging DRAM architectures in 3D (HMC, Wide I/O) for big data applications
  - o   Non-volatile logic and memory design
- **Power and Thermal Management**
  - o   Power/thermal and reliability issue in 3D architecture
  - o   Power management in emerging non-volatile memories
  - o   Power and energy optimization in VLSI circuits
  - o   Thermal management in emerging technologies such as 3D
  - o   Reliability-aware memory design
  - o   Dynamic power/thermal management in multi/many-core systems
  - o   Energy efficiency and power management in enterprise datacenter

PUBLICATIONS

Publication Summary:

Topic	Number of Publications	Conference/Journal
Computer Security	14	CHES, DAC, TODAES, ASPDAC, IOLTS, ISVLSI, ISQED, ICCD, DATE, CASES, ICCAD
Applied Machine Learning	23	ASPDAC, CASES, DAC, DATE, ICCD, ISQED, IISWC, FCCM, ICCAD, CODES-ISSS, CCGRID, ASAP, GLSVLSI, ASPDAC, ICCD, TVLSI
Big Data Computing	26	CCGRID, SoCC, TSMSCS, TOMPECS, TECS, JPDC, JPDC, TMSCS, ICCD, DAC, IISWC, ISPASS, FCCM, Big Data, CF, SAMOS, IGSC, DATE
Resource Management	56	ISCA, HPCA, TODAES, TVLSI, JETC, TVLSI, TVLSI, MR, TECS, CAL, JSC, TVLSI, LCTES, SAMOS, IGSC, GLSVLSI, ICCD, DATE, DAC

Conference/Journal	Number of Publications	Conference/Journal	Number of Publications	Conference/Journal	Number of Publications
DAC	7	ISCA	1	CASES	5
DATE	7	HPCA	1	CODES-ISSS	2
ISLPED	3	ISPASS	1	IISWC	3
ASPDAC	4	TVLSI	9	ISQED	7
ICCD	11	GLSVLSI	8	TMSCS	2
TODAES	2	TECS	3	SoCC	1
TOMPECS	1	JETC	2	CAL	1
CHES	1	LCTES	1	ISVLSI	3
ICCAD	1	CF	3	ISCAS	3
FCCM	3	Big Data	2	CCGRID	2

Journal Papers

(24) *"Programmable Gates Using Hybrid CMOS-STT Design to Prevent IC Reverse Engineering"*
Theodore Winograd, Hassan Salmani, Hamid Mahmoodi, Kris Gaj, **Houman Homayoun**
ACM Transactions on Design Automation of Electronic Systems, Special issue on Internet of Things
System Performance, Reliability, and Security, 2018 (Accepted)   ***TODAES***

(23) *"Optimal Allocation of Computation and Communication in an IoT Network"*
Abhimanyu Chopra, Hakan Aydin, Setareh Rafatirad, **Houman Homayoun**
ACM Transactions on Design Automation of Electronic Systems, Special issue on Internet of Things
System Performance, Reliability, and Security, 2018 (Accepted)   ***TODAES***

(22) *"Hardware Accelerated Mappers for Hadoop MapReduce Streaming"*
Katayoun Neshatpour, Maria Malik, **Houman Homayoun**
IEEE Transactions on Multi-Scale Computing Systems, 2018 (Accepted)   ***TMSCS***

(21) *"System and Architecture Level Characterization of Big Data Applications on Big and Little Core
Server Architectures"*
Maria Malik, Katayoun Neshatpour, Setareh Rafatirad, **Houman Homayoun**
ACM Transactions on Modeling and Performance Evaluation of Computing Systems, 2018
(Accepted)   ***TOMPECS***

(20) *"Low Overhead CS-based Heterogeneous Framework for Big Data Acceleration"*
Amey Kulkarni, Colin Shea, Tahmid Abtahi, **Houman Homayoun** and Tinoosh Mohsenin
ACM Transaction on Embedded Computing Systems, 2018 (Accepted)   ***TECS***

(19) *"Specialized Hardware-Supported Malware Detection Using Machine Learning Techniques"*
Hossein Sayadi, Nisarg Patel and **Houman Homayoun**
IEEE Transactions on Dependable and Secure Computing, 2019 (under review).   ***TDSC***

(18) *"Big vs Little Core for Energy-Efficient Hadoop Computing"*
Maria Malik; Katayoun Neshatpour; Setareh Rafatirad; Rajiv V Joshi; **Houman Homayoun**
Elsevier Journal of Parallel and Distributed Computing, Special Issue on Systems for Learning,
Inferencing, and Discovering (SLID), 2017 (Accepted)   ***JPDC***

(17) *"Smart Grid on Chip: Work Load-Balanced On-Chip Power Delivery"*
Divya Patahk, **Houman Homayoun**, Ioannis Savidis
IEEE Transactions on Very Large Scale Integration (VLSI) Systems, DOI: 10.1109/
TVLSI.2017.2699644, 2017   ***TVLSI***

(16) *"Energy-Efficient Acceleration of MapReduce Applications Using FPGAs"*
Katayoun Neshatpour; Maria Malik; **Houman Homayoun**
Elsevier Journal of Parallel and Distributed Computing, Special Issue on Systems for Learning,
Inferencing, and Discovering (SLID) (Accepted)   ***JPDC***

(15) "*An Energy Efficient Programmable Manycore Accelerator for Personalized Biomedical
Applications*"
Adam Page, Adwaya Kulkarni, Nasrin Attaran, Ali Jafari, Maria Malik, **Houman Homayoun**, and
Tinoosh Mohsenin
IEEE Transactions on Very Large Scale Integration (VLSI) Systems (Accepted)   ***TVLSI***

(14) "*Heterogeneous HMC+DDRx Memory Management for Performance-Temperature Trade-offs*"
Mohammad Hossein Hajkazemi, Mohammad Khavari Tavana, Tinoosh Mohsein, **Houman
Homayoun**
ACM Journal on Emerging Technologies in Computing, 2017. (Accepted)   ***JETC***

(13) "*Sparse Regression Driven Mixture Important Sampling for Memory Design*"
Maria Malik, Rajiv Joshi, Rouwaida Kanj, Shupeng Sun, **Houman Homayoun**, Tong Li
IEEE Transactions on Very Large Scale Integration (VLSI) Systems (Accepted)   ***TVLSI***

(12) "*Hadoop Workloads Characterization for Performance and Energy Efficiency Optimizations on*   ***TMSCS***

Microservers"
Maria Malik, Katayoun Neshatpour, Avesta Sasan, Setareh Rafatirad, **Houman Homayoun**
IEEE Transactions on Multi-Scale Computing Systems (Accepted)

(11) "*ElasticCore: A Dynamic Heterogeneous Platform with Joint Core and Voltage/Frequency Scaling*"   ***TVLSI***
Mohammad Khavari Tavana, Mohammad Hossein Hajkazemi, Divya Pathak, Ioannis Savidis, **Houman Homayoun**
IEEE Transactions on Very Large Scale Integration (VLSI) Systems (Accepted)

(10) "*Enhancing Power, Performance, and Energy-efficiency in Chip Multiprocessors Exploiting Inverse*   ***TVLSI***
Thermal Dependence"
Katayoun Neshatpour, Wane Burleson, Amin Khajeh, **Houman Homayoun**
IEEE Transactions on Very Large Scale Integration (VLSI) Systems (Accepted)

(9) "*Reliability analysis of spin transfer torque based look up tables under process variations and NBTI*   ***MR***
aging".
Ragh Kuttappa, **Houman Homayoun**, Hassan Salmani, Hamid Mahmoodi.
Elsevier Microelectronics Reliability Journal, Volume 62, p 156-166, 2016.

(8) "*Using a Flexible Fault-Tolerant Cache to Improve Reliability for Ultra Low Voltage Operation*".   ***TECS***
Abbas Banaiyanmofrad, **Houman Homayoun**, Nikil Dutt.
ACM Transactions on Embedded Computing. 14, no. 2 (2015): 32.

(7) "*Resistive Computation: A Critique*".   ***CAL***
Hamid Mahmoodi, Sridevi Srinivasan Lakshmipuram, Manish Arora, Yashar Asgarieh, **Houman Homayoun**, Bill Lin and Dean M.Tullsen.
IEEE Computer Architecture Letters, DOI 10.1109/L-CA.2013.23, 2014.

(6) "*Multi-Copy Cache: A Highly Energy Efficient Cache Architecture*".   ***TECS***
Arup Chakraborty, **Houman Homayoun**, Amin Khejah, Nikil Dutt, Ahmed Eltawil, Fadi Kurdahi.
ACM Transactions on Embedded Computing Systems (TECS), 2014.

(5) "*Variation Trained Drowsy Cache (VTD-Cache): A History Trained Variation Aware Drowsy*   ***TVLSI***
 Cache for Fine Grain Voltage Scaling".
Avesta Makhzan, Kiarash Amiri, **Houman Homayoun**, Ahmed Eltawil, Fadi J. Kurdahi.
IEEE Transactions on Very Large Scale Integration (VLSI) Systems, 2010 (TVLSI). VOL. 20, Issue 4, pp: 630-642. April 2012.

(4) "*MZZ-HVS: Multi Modes Zig-Zag Horizontal and Vertical Sleep Transistor Sharing to Reduce*   ***TVLSI***
 Leakage Power in On-Chip SRAM Peripheral Circuits".
Houman Homayoun, Avesta Sasan, Alex Veidenbaum, Hsin-Cheng Yao, Shahin Golshan, Payam Heydari.
IEEE Transactions on Very Large Scale Integration (VLSI) Systems, (TVLSI), VOL. 19, NO. 12, DECEMBER 2011.

(3) "*On Leakage Power Optimization in Clock Tree Networks for ASICs and General-Purpose*
Processors"   ***JSC***
Houman Homayoun, Shahin Golshan, Eli Bozorgzadeh, Alex Veidenbaum, Fadi Kurdahi.
Elsevier Journal of Sustainable Computing, Volume 1, Issue 1, March 2011, Pages 75-87 **(Invited paper)**.

(2) "*Inquisitive Defect Cache: A Means of Combating Manufacturing Induced Process Variation*".
Avesta Makhzan, **Houman Homayoun**, Ahmed Eltawil, Fadi J. Kurdahi.   ***TVLSI***
IEEE Transactions on Very Large Scale Integration (VLSI) Systems, 2010 (TVLSI), VOL. 19, NO. 9, SEPTEMBER 2011.

(1) "*Reducing Power in All Major CAM and SRAM Based Processor Units via Centralized, Dynamic*
Resource Size Management".
Houman Homayoun, Avesta Sasan, Alex Veidenbaum, Jean-Luc Gaudiot.   ***TVLSI***
IEEE Transactions on Very Large Scale Integration (VLSI) Systems, 2010 (TVLSI), VOL. 19, NO. 11, NOVEMBER 2011.

DOCKET ALARM

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts

Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research

With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips

Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS
Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS
Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS
Sync your system to PACER to automate legal marketing.

fastcase®
Smarter legal research.