

Mohamed Zahran

mzahran@acm.org

<http://www.mzahran.com>

- EDUCATION**
- ◇ **University of Maryland**, College Park, MD.
Ph.D. in Electrical and Computer Eng., May 2003
Dissertation Title: “Hierarchical Multithreading Architecture”
Adviser: Prof. Manoj Franklin GPA = 3.9/4.0
 - ◇ **Cairo University**, Egypt.
M.Sc. in Computer Engineering, July 1999.
Thesis title: “A New Adaptive Genetic Algorithm”
 - ◇ **Cairo University**, Egypt.
B.Sc. in Computer Engineering, July 1997.
Graduation Project: “Computer Chess”
- RESEARCH INTERESTS**
- ◇ **AI Support for Architecture**
 - ◇ **Graphics Processing Units and other Hardware Accelerators**
 - ◇ **Memory hierarchy design for Heterogeneous Computing**
 - ◇ **Hardware-software interaction**
- WORK EXPERIENCE**
- ◇ **New York University** (Sep 2018 - present)
Professor (clinical) with Computer Science Department, Courant Institute of Mathematical Sciences, NYU
 - ◇ **New York University** (Sep 2012 - Aug 2018)
Associate Professor (clinical) with Computer Science Department, Courant Institute of Mathematical Sciences, NYU
 - ◇ **New York University** (Jan 2010 - May 2012)
Adjunct Associate Professor with Computer Science Department, Courant Institute of Mathematical Sciences, NYU
 - ◇ **Polytechnic Institute of New York University** (Jan 2010 - Dec 2011)
Research Associate Professor, Electrical and Computer Engineering Department
 - ◇ **Consultancy Work** (2007-2010)
 - Polytechnic Institute of NYU:
 - Designing power-efficient supercomputer based on hybrid FPGA and general purpose processors using switched networks.
 - Architecture support for Trusted Platform Management for single core, multicore, and manycore processors.
 - Cairo Microsoft Innovation Center (CMIC):
Proposed and supervised research projects related to high-performance computing. The projects I have contributed to are related to parallelization and analysis of bioinformatics applications.
 - ◇ **City University of New York - City College** (Sep 04 - May 10)
Department of Electrical Engineering and the Computer Engineering program,
Assistant Professor (Sep 04 - Aug 09), then Adjunct Associate Professor (Sep 09 - May 10)
 - ◇ **The George Washington University** (Sep 03 - Jul 04)
Research Scientist with the Department of Electrical and Computer Engineering.
Duties included directing two groups of graduate students:

- The first group did research on the specifications and performance of a parallel version of the C language, called unified parallel C (UPC).
 - The second group focused on reconfigurable architecture, and how to dynamically reconfigure a device to adapt to program(s) behavior.
- ◇ **University of Maryland at College Park**
- Research Assistant (Aug 99 - Aug 03):
Department of Electrical and Computer Engineering
Research topic: The Microarchitecture of Speculative Multithreaded Processors.
 - Teaching Assistant (Aug 99 - May 2001):
Department of Electrical and Computer Engineering
Assisted in teaching: Digital Logic Design, Computer Organization.
Duties: holding weekly recitation classes, office hours as well as grading homework assignments and exams.
- ◇ **Cairo University, Egypt** (Sep 97 - May 99)
Assistant Lecturer:
Teaching: C programming language, Digital Logic Design,
Computer Architecture, Software Engineering and Operating Systems.
- ◇ **American University in Cairo (AUC)** (Spring 98)
Lecturer:
Center of Adult and Continuing Education (CACE),
Duties: Teaching computer architecture and assembly language programming course.
- ◇ **IBM Egypt** (Summer 96)
Summer Intern:
Performance analysis and tuning of IBM RISC/6000.
- ◇ **AT&T Global Information System, Egypt**
Summer Intern:
- (Summer 95): UNIX operating system and Unix Shell programming workshops, Data communication and LAN and WAN concepts workshops.
 - (Summer 94): Windows applications workshop, C programming language and Clipper application workshops.

HONORS AND AWARDS ◇ Elected IEEE Distinguished Contributor (2021)

- ◇ ACM Distinguished Speaker (2019 - 2025)
- ◇ Best poster award for CS and Mathematics, for my student Daniel Cohen on his research work with me, in the 41st Undergraduate Research Conference, College of Arts and Science, NYU, 2015.
 - ◇ First place in ACM Student Research Competition with my undergraduate student Yi (Louisa) Lu, held in conjunction with The International Symposium on Code Generation and Optimization (CGO) 2014.
 - ◇ Best Poster Award in Trusted Infrastructure Workshop (TIW), June 2010.
 - ◇ First place in ACM Student Research Competition with my Ph.D student Bushra Ahsan, held in the 18th International Conference on Parallel Architectures and Compilation Techniques (PACT) 2009.

- ◇ Elected Senior member of ACM.
- ◇ Elected Senior member of IEEE.
- ◇ Certificate of Recognition at City University of New York on the occasion of the “Salute to Scholars” event in Dec 2008 in honor of outstanding scholarly achievements.
- ◇ Top three nationwide (second place) in ACM Student Research Competition with my undergraduate student Jerry Backer, 2007.
- ◇ Awarded grant from IBM to attend 39th International Symposium on Microarchitecture (MICRO), 2006.
- ◇ Awarded grant from ACM SIGARCH for young faculties to attend International Conference on Microarchitecture, held in Oregon, December 2004.
- ◇ Best paper award in International Conference on Computer Design (ICCD), San Jose, CA, October, 2003.
- ◇ ACM/SIGARCH conference student grant award: ISCA/FCRC 2003
- ◇ Best Graduate student scientific talk in Electrical and Computer Engineering dept., University of Maryland, fall 2002.
- ◇ Honors with Distinction (First Place) in University of Maryland Graduate Research Interaction Day (GRID), 2002.
- ◇ Honors with Distinction at M.Sc (top 5), Cairo University, 1999.
- ◇ Honors with Distinction at B.Sc, Cairo University 1997.

PUBLICATIONS **Books:**

- **Mohamed Zahran**, *Heterogeneous Computing: Hardware and Software Perspectives*, ACM Books, ISBN: 978-1-4503-6097-5, 2019.

Book Chapters:

- Bushra Ahsan and **Mohamed Zahran**, *Cache Hierarchy for Manycore Processors*, in the book *Multicore Computing: Algorithms, Architectures, and Applications*, Chapman and Hall/CRC, Dec 2013.

Refereed Papers:

1. D. Tantawy, **M. Zahran** and A. G. Wassal, *PTcomp: Post-Training Compression Technique for Generative Adversarial Networks*, in IEEE Access, vol. 11, pp. 9763-9774, 2023.
2. Dina Tantawy, **Mohamed Zahran**, Amr Wassal, *A Survey on GAN Acceleration Using Memory Compression Technique*, Journal of Engineering and Applied Science, 68:47, Springer Nature, 2021.
3. Nick Greenquist, Doruk Kilitcioglu, **Mohamed Zahran** and Anasse Bari, *GPU Accelerated Matrix Factorization for Recommender Systems*, the 6th IEEE International Conference on Big Data Analytics (ICBDA 2021), March 2021.
4. Antonio Mallia, Michal Siedlaczek, Torsten Suel, and **Mohamed Zahran**, *GPU-Accelerated Decoding of Integer Lists*, in The 28th ACM International Conference on Information and Knowledge Management (CIKM), Beijing, China, November 2019.
5. Tulsi Jain, Nitish Agarwal, and **Mohamed Zahran**, *Performance Prediction for Multithreaded Applications*, in The 2nd International Workshop on AI-assisted Design for Architecture (AIDArc), to held in conjunction with the International Symposium on Computer Architecture, June 2019.
6. **Mohamed Zahran** and Marsha Berger, *Parallel Computing At The Undergraduate Level: Lessons Learned and Insights*, in Workshop on Computer Architecture Education Held in conjunction with 46th International Symposium on Computer Architecture, June 2019.

7. Mahmoud Khairy, Amr Wassal, and **Mohamed Zahran**, *A survey of architectural approaches for improving GPGPU performance, programmability and heterogeneity*, Elsevier Journal of Parallel and Distributed Computing, Volume 127, May 2019, Pages 65-88.
8. Chris Quackenbush and **Mohamed Zahran**, *Beyond Profiling*, in The 1st International Workshop on AI-assisted Design for Architecture (AIDArc), to held in conjunction with the International Symposium on Computer Architecture, June 2018.
9. Mahmoud Khairy , **Mohamed Zahran**, and Amr Wassal, *SACAT: Streaming-Aware Conflict-Avoiding Thrashing-Resistant GPGPU Cache Management Scheme*, IEEE Transactions on Parallel and Distributed Systems, vol 28, issue 6, June 2017.
10. Numair Khan and **Mohamed Zahran**, *Space-efficient Pointwise Computation of the Distance Transform on GPUs*, in 7th IEEE Workshop Parallel / Distributed Computing and Optimization (PDCO 2017), in conjunction with 31st IEEE International Parallel & Distributed Processing Symposium (IPDPS), May 2017.
11. Chris Rohlfs and **Mohamed Zahran**, *Optimal Bandwidth Selection for Kernel Regression Using a Fast Grid Search and a GPU*, in 7th IEEE Workshop Parallel / Distributed Computing and Optimization (PDCO 2017), in conjunction with 31st IEEE International Parallel & Distributed Processing Symposium (IPDPS), May 2017.
12. **M. Zahran**, *Heterogeneous Computing: Here to Stay*, ACM Queue magazine, Nov/Dec 2016, also appears in Communications of the ACM magazine, March 2017.
13. **Mohamed Zahran**, *Brain-Inspired Machines What, Exactly, Are We Looking for?*, IEEE Pulse Magazine, March 2016.
14. Mahmoud Khairy, **Mohamed Zahran**, and Amr G. Wassal, *Efficient utilization of GPGPU cache hierarchy*, in the 8th Workshop on General-purpose processing using GPUs held in conjunction with the 20th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, 2015.
15. J. Rajendran, A. K. Kanuparthi, **M. Zahran**, S. Addepalli, G. Ormazabal, and R. Karri, *Securing processors against insider attacks: a circuit-microarchitecture co-design approach*, IEEE Design and Test of Computers Magazine, Vol 30, issue 2, Mar/Apr, 2013.
16. Arun K. Kanuparthi, **Mohamed Zahran**, and Ramesh Karri, *Architecture Support for Dynamic Integrity Checking*, IEEE Transactions on Information Forensics and Security, Vol. 7, Issue 1, pp. 321-332, 2012.
17. H. Chtioui, S. Niar Lamih, R. Ben-Atitallah, **M. Zahran**, Jl. Dekeyser, and M. Abid, *A Dynamic Hybrid Cache Coherency Protocol for Shared-Memory MPSoC Architectures*, International Journal of Computer Applications, Volume 47, Number 3, 2012.
18. Corey Malone, **Mohamed Zahran**, and Ramesh Karri, *Are Hardware Performance Counters a Cost Effective Way for Integrity Checking of Programs?*, The Sixth ACM Workshop on Scalable Trusted Computing, October 2011.
19. Artem Durytskyy, **Mohamed Zahran**, and Ramesh Karri, *Improving Robustness of GPUs by Making Use of Faulty Parts*, Proc. International Conference on Computer Design (ICCD11), October 2011.
20. Mohamed Salah Souahi, Smail Niar, **Mohamed Zahran**, Mohamed Benmohamed, *Towards Dynamic Cache Block Placement for Multi-processor NUCA*, IEEE International Conference on Microelectronics, December 2011.
21. Arun K. Kanuparthi, **Mohamed Zahran**, and Ramesh Karri, *Feasibility Study of Dynamic Trusted Platform Module*, Proc. International Conference on Computer Design (ICCD10), October 2010.
22. Ahmed Youssef, **Mohamed Zahran**, Mohab Anis, and Mohamed Elmasry, *On the Power Management of Simultaneous Multithreading Processors*, IEEE Transactions on VLSI Systems, pp. 1243-1248, Vol. 18, August 2010.

23. **Mohamed Zahran** and Sally A. McKee, *Global Management of Cache Hierarchies*, The ACM International Conference on Computing Frontiers (CF'10), Italy, May 2010.
24. Yufu Zhang, Ankur Srivastava, and **Mohamed Zahran**, *On-Chip Sensor Driven Efficient Thermal Profile Estimation Algorithms*, ACM Transactions on Design Automation of Electronic Systems, Volume 15 , Issue 3, May 2010.
25. Najla Alfaraj, H. Jonathan Chao, and **Mohamed Zahran**, *NBC: Network-based Cache Coherence Protocol for Multistage NoCs*, in The International SoC Design Conference (ISOC), 2009.
26. Kim Hazelwood and **Mohamed Zahran**, *Challenges and Opportunities at All Levels: Interactions Among Operating Systems, Compilers, and Multicore Processors*, ACM SIGOPS Operating System Review. Volume 43, Issue 2. April 2009.
27. Bushra Ahsan and **Mohamed Zahran**, *Managing Off-Chip Bandwidth: A Case for Bandwidth-Friendly Replacement Policy*, in The 2nd Workshop on Managed Multi-Core Systems (MMCS'09), held in conjunction with ASPLOS 2009.
28. **Mohamed Zahran** and Sally A. McKee, *Adaptive Block Placement Policy for Cache Hierarchies* , in 3rd Workshop on statistical and Machine learning approaches to Architectures and compilaTion (SMART'09), held in conjunction with HiPEAC'09.
29. Bushra Ahsan and **Mohamed Zahran**, *Cache Performance, System Performance, and Off-Chip Bandwidth... Pick any Two*, in 3rd workshop Interconnection Network Architectures: On-Chip, Multi-Chip (INA-OCMC), held in conjunction with HiPEAC 2009.
30. Yufu Zhang, Ankur Srivastava and **Mohamed Zahran**, *Chip Level Thermal Profile Estimation Using On-chip Temperature Sensors*, Proc. International Conference on Computer Design (ICCD08), October 2008.
31. **Mohamed Zahran**, *Cache Replacement Policy Revisited*, in The Annual Duplicating, Deconstructing, and Debunking (WDDD)held in conjunction with the International Symposium on Computer Architecture (ISCA), 2007.
32. **Mohamed Zahran**, Kursad Albayraktaroglu, and Manoj Franklin, *Non-Inclusion Property in Multi-Level Caches Revisited*, in the International Journal of Computers and Their Applications, Special Issue on Techniques and Architectures for High Performance and Energy Efficient Computing Systems, Vol 14, Num 2, June 2007.
33. **Mohamed Zahran**, *Cache Hierarchy for 100 On-Chip Cores*, Fifth Annual Boston Area Architecture (BARC), Jan 2007.
34. **Mohamed Zahran**, *On Cache Blocks Behavior*, in International Computer Engineering Conference (ICENCO), Dec 2006.
35. **Mohamed Zahran** and Manoj Franklin, *RHT: A Context-Based Return Address Predictor*, in The 2006 International Conference on Computer Design (CDES'06), Las Vegas, June 2006.
36. **Mohamed Zahran** and Anasua Bhowmik, *Bandwidth-Friendly Cache Hierarchy*, in The 2006 International Conference on Computer Design (CDES'06), Las Vegas, June 2006.
37. **Mohamed Zahran** and Anasua Bhowmik, *Hybrid Compiler and Microarchitecture Technique for Cache Traffic Optimization*, in Interaction between Compilers and Computer Architectures (INTERACT 9), Feb 2005.
38. Francois Cantonnet, Yiyi Yao, **Mohamed Zahran** and Tarek El-Ghazawi, *Productivity Analysis of the UPC Language*, in 3rd International Workshop on Performance Modeling, Evaluation, and Optimization of Parallel and Distributed Systems (PMEO-PDS), to be held in conjunction with the International Parallel and Distributed Processing Symposium (IPDPS 2004).

39. **Mohamed Zahran** and Manoj Franklin, *Dynamic Thread Resizing for Speculative Multithreaded Processors*, in International Conference on Computer Design (ICCD), San Jose, CA, October, 2003. (**BEST PAPER AWARD**)
40. **Mohamed Zahran**, Manoj Franklin and Renju Thomas, *Confidence Estimation for Register Value Communication in Speculative Multithreaded Architectures*, in first value prediction workshop (VPW1), held in conjunction with the 30th Annual International Symposium on Computer Architecture (ISCA), San Diego, California, 2003.
41. **Mohamed Zahran**, *On Cache Memory Hierarchy for Chip-Multiprocessor*, in MEDEA workshop held in conjunction with PACT 2002 Conference, Charlottesville, Virginia, 2002. Also Appeared in ACM Computer Architecture News, Vol 31, No. 1, March 2003.
42. **Mohamed Zahran** and Manoj Franklin, *Return Address Prediction in Speculative Multithreaded Environments*, in Int'l Conference on Hi-Performance Computing (HiPC'02), Bangalore, India, 2002.
43. **Mohamed Zahran** and Manoj Franklin, *A Feasibility Study of Hierarchical Multi-threading*, in International Parallel and Distributed Processing Symposium (IPDPS 2002), Marriott Marina, Fort Lauderdale, Florida, 2002.
44. **Mohamed Zahran** and Manoj Franklin, *Hierarchical Multi-threading For Exploiting Parallelism at Multiple Granularities*, Workshop on MULTITHREADED EXECUTION, ARCHITECTURE and COMPILATION (MTEAC-5), Austin, Texas, 2001.
45. **Mohamed Zahran**, Ashraf Abdel-Wahab and Samir Shaheen, *Adaptive Genetic Algorithm for Multiprocessor Scheduling*, poster presentation at the Genetic and Evolutionary Computation Conference (GECCO), Orlando, 1999.

Refereed Posters, Abstracts, Invited Papers, and Technical Reports:

1. Marcelo Pias, Brett Becker, Qiao Xiang, **Mohamed Zahran**, and Monica Anderson, *Should Quantum Processor Design be Considered a Topic in Computer Architecture Education?* In Proceedings of the 53rd ACM Technical Symposium on Computer Science Education V. 2 (SIGCSE 2022).
2. **Mohamed Zahran**, *The Future of High-Performance Computing*, 17th International Computer Engineering Conference (ICENCO), pp 129 - 134, 2021 (Invited Paper).
3. **M. Zahran**, *Multicore processors: Status quo and future directions*, in 10th International Computer Engineering Conference (ICENCO), Dec 2014 (Invited Paper).
4. Yi (Louisa) Lu and **Mohamed Zahran**, *Unleashing the Power of General Purpose Graphics Processing Units*, poster at ACM Students Research Competition, held in Conjunction with The International Symposium on Code Generation and Optimization, FL, 2014 (**Winner of ACM SRC, First Prize**).
5. Arun K. Kanuparthi, **Mohamed Zahran**, and Ramesh Karri, *On-Chip Dynamic Trusted Platform Module*, Trusted Infrastructure Workshop (TIW), June 2010, Pittsburgh, PA, (**Best Poster Award**).
6. Bushra Ahasan and **Mohamed Zahran**, *A Hybrid Compiler-Architecture Technique to Manage Off-Chip Traffic for Multicore Chips*, poster at ACM Students Research Competition, held in Conjunction with The 18th International Conference on Parallel Architecture and Compilation Techniques, NC, 2009 (**Winner of ACM SRC, First Prize**).
7. **Mohamed Zahran** and Sally A. McKee, *Enterprise-Like Cache Hierarchy Management in the Manycore Era*, position abstract, ACM SIGPLAN Workshop on Memory Systems Performance and Correctness (MSPC) held in conjunction with ASPLOS, March, 2008.

8. Bushra Ahsan, Fatma Omara and **Mohamed Zahran**, *Chip Multiprocessor: Challenges and Opportunities*, INFOS 2008. (Invited Paper)
9. Jerry Backer and **Mohamed Zahran**, *The Effect of Memory Bandwidth on Processor Performance*, poster at Richard Tapia Conference, FL 2007 (**Winner of ACM SRC Grand Final**).
10. Bushra Ahasan and **Mohamed Zahran**, *Cache Improvement Techniques Reconsidered: A Write-Buffer Case Study*, poster at Richard Tapia Conference, FL 2007.
11. A. Bhowmik and **M. Zahran**, *Cache Traffic Optimization*, Computer Science and Automation, Indian Institute of Science, India, 2005, IISc-CSA-TR-2005-1.
12. **Mohamed Zahran**, *On Cache Memory Hierarchy for Chip-Multiprocessor*, ACM Computer Architecture News, Vol 31, No. 1, March 2003.

PRESENTATIONS
& TALKS

1. *Artificial Intelligence Leaps in Digital Banking*, University of Cranfield, UK, distinguished speaker, May 2023.
2. *Artificial Intelligence Leaps and Financial Applications*, University of Cranfield, UK, distinguished speaker, November 2021.
3. *AI Support for Computer Architecture*, SEMINAR IN ADVANCES IN COMPUTING, University of South Carolina (UofSC), September 2021.
4. *AI in the Middle East and North Africa: visions and realities*, panel organized by AI for Good and UN, June 2021.
5. *AI Support for Architecture: Challenges and Opportunities*, AMD Research ,July 2020.
6. *Toward Exascale Machine: Challenges and Opportunities*, IBM T. J. Watson lab , April 2017.
7. *Architecture Support for Big Data*, Bloomberg, November 2016.
8. Panel at IBM Research Workshop on Architectures for Cognitive Computing and Datacenters, IBM T. J. Watson lab , October 2016.
9. *Heterogeneous Computing: Hardware and Software Perspective*, ACM Applicative, June 2016.
10. *Hardware Advances for the Big Data Era*, Center for Data Sciences, New York University, November 2015.
11. *Off-Chip Bandwidth: The New Wall in The Multicore Era*, Computer and Information Sciences Departmental seminar series, University of Delaware, March 2009.
12. *Multicore Chips and The Green Revolution*, IT Symposium by Uptime Institute, April 2009.
13. *Off-Chip Bandwidth: The New Wall in The Multicore Era*, Intel-VSSAD, Hudson, MA, May 2009.
14. *Adaptive Block Placement Policy for Cache Hierarchies* in SMART'09, held in conjunction with HiPEAC'09.
15. *Cache Performance, System Performance, and Off-Chip Bandwidth... Pick any Two*, INA-OCMC, held in conjunction with HiPEAC'09.
16. *Enterprise-Like Cache Hierarchy Management in the Manycore Era*, ACM SIGPLAN Workshop on Memory Systems Performance and Correctness (MSPC) held in conjunction with ASPLOS, March, 2008.
17. *Attacking The Von-Neumann Bottleneck: Cache Hierarchy in The Chip Multiprocessor Era*, as part of the departmental seminar series of Fall 2007, the ECE/CS dept of Polytechnic University.

18. *Attacking The Von-Neumann Bottleneck: Cache Hierarchy in The Chip Multiprocessor Era*, as part of the departmental seminar series of Fall 2007, the CS dept, university of Massachusetts at Amherst.
19. *Cache Replacement Policy Revisited*, The Annual Workshop on Duplicating, Deconstructing, and Debunking (WDDD) held in conjunction with the International Symposium on Computer Architecture (ISCA), 2007.
20. *Attacking the Von-Neumann Bottleneck: Smart and Scalable Cache Hierarchy in The Chip Multiprocessor Era*, IBM T. J. Watson, Feb 2007.
21. *Computer Hardware: A Different Perspective*, keynote speech, 2nd International Computer Engineering Conference (ICENCO), Dec 2006.
22. *RHT: A Context-Based Return Address Predictor*, The 2006 International Conference on Computer Design (CDES'06), Las Vegas, June 2006.
23. *Bandwidth-Friendly Cache Hierarchy*, The 2006 International Conference on Computer Design (CDES'06), Las Vegas, June 2006.
24. *Chip Multithreading: Issues and Challenges*, part of the departmental seminar series of fall 2005 in the ECE dept, university of Massachusetts at Amherst.
25. *Cache Traffic Optimization and Cache on Demand System*, at a meeting at NSF with CSR program directory Prof. Peter Varman, June 2005.
26. *Hybrid Compiler and Microarchitecture Technique for Cache Traffic Optimization*, INTERACT-9 (Interaction between Compilers and Computer Architectures), Feb 2005.
27. *Confidence Estimation for Register Value Communication in Speculative Multithreaded Architectures*, first value prediction workshop (VPW1), held in conjunction with the 30th Annual International Symposium on Computer Architecture (ISCA), San Diego, California, 2003.
28. *Speculative Multithreading...The Future of Microprocessors*, Electrical and Computer Engineering Graduate Students Association (ECEGSA) Graduate Student Seminar, Fall 2002, **BEST SCIENTIFIC TALK AWARD**
29. *Microprocessors...Can We Make Further Progress*, University of Maryland Graduate Research Interaction Day (GRID), Spring 2002, **BEST TALK AWARD**
30. *On Cache Memory Hierarchy for Chip-Multiprocessor*, MEDEA workshop, held in conjunction with PACT, Virginia, 2002.
31. *Feasibility Study of Hierarchical Multithreading*, IPDPS Conference, Florida, 2002.

GRANTS

◇ External Funding:

- Equipment: Xilinx: 2017, several FPGA boards and software tools to support our research and teaching in heterogeneous computing.
- CRA-CREU 2015/2016, Computing Research Association, title: "Dealing With Memory in the Parallel Computing Era", (\$7,500).
- CRA-CREU 2013/2014, Computing Research Association, title: "General Purpose Graphics Processor Units (GPGPUs) Bottleneck Analysis", (\$4,500).
- NSF-SGER 2007/2008: "SGER: Exploring the Potential for Software-Informed Hardware Reconfigurability in the Memory Hierarchy of Embedded Systems" (\$ 40,000).
- NSF-CBET 2007-2010, "MRI: Acquisition of an advanced micro-Computed Tomography imaging facility", (\$ 339,450) [Co-PI].
- CRA-CREU 2007/2008, Computing Research Association, title: "Toward Better Memory Hierarchy for Chip-Multiprocessor", (\$9,500).
- CRA-CREU 2006/2007, Computing Research Association, title: "Effect of Bus Traffic on Cache Hierarchy Performance", (\$3,500).

◇ **Internal Funding:**

- NYU - University Research Challenge Fund,(2014/2015), *Generalized Profiling: Extracting Wisdom from Profiling Data*, (\$15,000).
- PSC-CUNY grant 40 (2009/2010), The Professional Staff Congress-City University of New York award for research on *Off-Chip Bandwidth Management for Multicore Processors*,(\$2,700).
- PSC-CUNY grant 38 (2007/2008), The Professional Staff Congress-City University of New York award for research on *Global Replacement Policy*,(\$3,950).
- PSC-CUNY grant 36 (2005/2006), The Professional Staff Congress-City University of New York award for research on *Improving Cache Memory Performance in Current and New Microarchitecture Environments*(\$2,320).

PROFESSIONAL
ACTIVITIES

◇ **General Chair:**

- The IEEE/ACM 49th International Symposium on Computer Architecture (ISCA 2022), co-Chair with Valentina Salapura

◇ **Review Panel:**

- National Science Foundation (NSF)
- Department of Energy (DoE)

◇ **Program committee member:**

- The 56th IEEE/ACM International Symposium on Microarchitecture (MICRO 2023)
- The 50th IEEE/ACM International Symposium on Computer Architecture (ISCA 2023)
- International Workshop on Performance Analysis of Machine Learning Systems (FAST-PATH 2023)
- The 40th IEEE International Conference on Computer Design (ICCD 2022)
- The 39th IEEE International Conference on Computer Design (ICCD 2021)
- The 38th IEEE International Conference on Computer Design (ICCD 2020)
- The ACM International Conference on Computing Frontiers (CF 2020)
- The 20th IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing (CCGRID 2020)
- The International Conference for High Performance Computing, Networking, Storage, and Analysis (SC 2019)
- The 33rd IEEE International Parallel & Distributed Processing Symposium (IPDPS 2019)
- The ACM International Conference on Computing Frontiers (CF 2019)
- The ACM International Conference on Computing Frontiers (CF 2018)
- The ACM International Conference on Computing Frontiers (CF 2017)
- The 23rd IEEE Symposium on High Performance Computer Architecture (HPCA 2017)
- The 30th IEEE International Parallel & Distributed Processing Symposium (IPDPS 2016)
- The ACM International Conference on Computing Frontiers (CF 2016)
- The ACM International Conference on Computing Frontiers (CF 2015)
- The 29th IEEE International Parallel & Distributed Processing Symposium (IPDPS 2015)
- The 11th IFIP International Conference on Network and Parallel Computing (NPC 2014)

- The 11th IEEE/IFIP International Conference on Embedded and Ubiquitous Computing (EUC-13)
- Student Research Symposium held in High-Performance Computing Conference (SRS-HiPC 2013)
- The 10th IEEE/IFIP International Conference on Embedded and Ubiquitous Computing (EUC 2012)
- The 25th International Conference on Supercomputing (ICS-2011)
- The 24th IEEE International Parallel and Distributed Processing Symposium (IPDPS-2010)
- The International Symposium on Code Generation and Optimization (CGO-2010)
- The Fourth international Workshop on Automatic Performance Tuning (iWAPT 2009)
- The 11th IEEE International Conference on High Performance Computing and Communications (HPCC-09)
- The 2nd International Conference on Computer Science and its Applications (CSA 2009)
- IEEE International Conference on Microelectronics (ICM) 2007, 2008, and 2009
- The International Computer Design Conference (CDES-2006)
- ◇ **Editorial:**
 - Steering Committee member, IEEE Transaction on Multi-Scale Computing (term 2018-2020).
 - Co-guest editor for a special issue of ACM Operating System Review about The interaction among Compilers, Operating Systems, and Multicore; April 2009.
- ◇ **Organizing Committee Member:**
 - Publicity chair for the ACM International Conference on Computing Frontiers 2014 (CF 2014)
 - Publication Chair for the ACM International Conference on Computing Frontiers 2011 and 2013
 - Workshops and tutorials co-chair for the 23rd International Conference on Supercomputing (ICS 2009)
 - Publicity chair for the Eighteenth International Conference on Parallel Architectures and Compilation Techniques (PACT 2009)
 - Co-organizer of 2nd Reconfigurable and Adaptive Architecture Workshop (RAAW), held in conjunction of The 40th Annual IEEE/ACM International Symposium on Microarchitecture, Dec 2007.
 - Co-organizer of 1st Reconfigurable and Adaptive Architecture Workshop (RAAW), held in conjunction of The 39th Annual IEEE/ACM International Symposium on Microarchitecture, Dec 2006.
- ◇ **Reviewing activities:**
 - Elsevier Microprocessors and Microsystem journal
 - The International Journal of Parallel Programming (IJPP)
 - International Journal of Parallel, Emergent, and Distributed Systems
 - Reviewer for Air Force Office of Scientific Research
 - IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems
 - ACM Transactions on Architecture and Code Optimization (TACO)
 - IEEE Computer Architecture Letters (CAL)

- Journal Of Circuits, Systems, and Computers
- Journal of Computers and Their Applications
- International journal of embedded systems (IJES)
- International Conference on Supercomputing (ICS)
- International Symposium on Computer Architecture (ISCA)
- International Symposium on High Performance Computer Architecture (HPCA)
- International Symposium on Circuits and Systems (ISCAS)
- International Parallel and Distributed Processing Symposium (IPDPS)
- International Conference on High Performance Computing (HiPC)
- IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)
- International Computer Engineering Conference (ICENCO).

- SOCIETIES AND AFFILIATIONS
- ◇ Full Member of Sigma Xi Honor society since 2002
 - ◇ Senior member of IEEE (including Computer Society).
 - ◇ Member of the following IEEE Computer society technical committees :
 - Computer Architecture (TCCA)
 - Microprocessors and MicroComputers (TCMCMP)
 - Microprogramming and Microarchitecture (TCMARCH)
 - ◇ Senior Member of ACM.
 - ◇ Member of the following ACM Special Interest Groups (SIG) :
 - Computer Architecture (SIGARCH)
 - Microprogramming (SIGMICRO)
 - Operating Systems (SIGOPS)
 - Computer Science Education (SIGCSE)
 - Design Automation (SIGDA)
 - Programming Languages (SIGPLAN)
 - ◇ Member of Egyptians Syndicate Professional Engineers society, since 1997.

- TEACHING EXPERIENCE
- ◇ **Multicore Processors: Architecture and Programming**
In Computer Science Department of New York University
A graduate course that studies multicore processors, including homogeneous and heterogeneous, from hardware and software perspective.
 - ◇ **Graphics Processing Units (GPUs): Architecture and Programming**
In Computer Science Department of New York University
A graduate course that studies GPUs from hardware and software perspective.
 - ◇ **Parallel Computing**
In Computer Science Department of New York University
An undergraduate course that introduces students to both the hardware and software aspects of parallel computing. Students are introduced to both shared-memory and distributed memory architectures. They learn to program in MPI (for distributed-memory systems), OpenMP (for shared-memory systems), and CUDA (for GPUs).

- ◇ **Computer Systems Organization**
In Computer Science Department of New York University
An undergraduate course that discusses what happens under the hood in a computer systems. It includes linkers, loaders, memory allocation, virtual memory, memory hierarchy, as well as C programming language and x86 assembly.
- ◇ **Virtual Machines: Concepts and Applications**
In Computer Science Department of New York University
A graduate course that studies the design and implementation of virtual machines, virtual machine monitors, as well other recent trends in virtualization. It discusses virtual machines across the disciplines that use them: operating systems, programming languages, and computer architecture.
- ◇ **Operating Systems**
In Computer Science Department of New York University
A course that discusses the main concepts of operating systems design.
I taught both the graduate and undergraduate editions.
- ◇ **Compilers Construction:**
In Computer Science Department of New York University
A graduate course in compilers design, that involves building a full compiler.
- ◇ **PAC II (Preparatory Accelerated Course):**
In Computer Science Department of New York University
A graduate course to prepare non-CS students for graduate studies. It explores the whole computing stack from algorithms to transistors.
- ◇ **Capstone Project:**
In Electrical Engineering Department at City College of City University of New York
Supervising several senior design projects, including: interfacing, software tool to help computer architecture researchers, and microprocessor design
- ◇ **Digital Computer Systems** (Spring 2005, Fall 2005, Spring 2006, fall 2006, Spring 2007, Fall 2007, Spring 2008, Fall 2008, Spring 2009, Fall 2009, Spring 2010) In Electrical Engineering Department at City College of City University of New York
This is a senior-level course about computer architecture.
- ◇ **Computer Engineering Laboratory:** (Fall 2004)
In Electrical Engineering Department at City College of City University of New York
This senior-level lab consists of several hands-on experiments in interfacing and FPGA.
I have also been faculty co-ordinator for that course (Spring 2005, Fall 2005, Spring 2006, and fall 2006)
- ◇ **Co-Teaching Graduate Course in Computer Architecture:** (Spring 04)
Advanced graduate course in computer architecture and microprocessor design
In ECE department at The George Washington University
Duties include helping in preparing the syllabus, exams, projects, and giving lectures.
- ◇ **Teaching assistant:** (Fall 99 - Spring 01)
In ECE department, University of Maryland at College Park. Assisted in teaching: Digital Logic Design, Computer Organization.
Duties: holding weekly recitation classes, office hours as well as grading homework assignments and exams.
- ◇ **Assistant Lecturer** (Sep 97 - May 99)
Cairo University, Egypt.
Teaching: C programming language, Digital Logic Design, Computer Architecture, Software Engineering and Operating Systems.
- ◇ **Lecturer** (Spring 98)
Center of Adult and Continuing Education (CACE),
American University in Cairo (AUC).

Duties: Teaching computer architecture and assembly language programming course.

ADVISING

◇ **Ph.D. Adviser of:**

- Bushra Ahsan, Dissertation Title: *Off-Bandwidth for Multicore Processors: The Next Big Wall*, First employment: postdoc at University of Cyprus in 2010/2011; Now at Intel Hudson, MA.

◇ **Ph.D committee member of the following students:**

- Ahmed Elkammar, City University of New York, August 2007
- Zahidur Rahman, City University of New York, June 2007
- Zhaoming Li, City University of New York, May 2007
- Qiang Song, City University of New York, May 2007
- Hassan Bajwa, City University of New York, Apr 2007
- Flavio De Angelis, City University of New York, 2005
- Kafi Hassan, City University of New York, 2005.
- Osama Hussein, City University of New York, 2005.
- Hooshang Sharif, City University of New York, 2005.

◇ **M.Sc. committee member of the following Students:**

- Hassan Bajwa, City College of City University of New York, Feb 2006
- Jonathan Cardenas, City College of City University of New York, Dec 2005
- Francois Cantonnet, ECE Dept, The George Washington University, Spring 04.

◇ **Supervisor of M.Sc. students for their M.Sc. thesis:**

- Kumar Prasun, *Configurable GPUs*, (Expected May 2023)
- Chris Quackenbush, *A New Method for Using Profiling Data to Improve Efficiency Across Multiple Applications*, Computer Science Department, Courant Institute of Mathematical Sciences, NYU (Graduated Spring 2016).
- Beihong Chen, *Cache Geometry and Architecture Acceleration for Machine Learning Approaches*, ECE Department, Tandon School of Engineering, NYU, (graduated summer 2016).
- Artem Durytskyy, ECE Department, Polytechnic Institute of NYU, (Graduated Spring 2011)
- Lakshmi Cuddalore Arivudainambi, City College of City University of New York, (Fall 2007)
- Preethi Gopinath, City College of City University of New York, (Fall 2007)
- Sravan Paruchuri, City College of City University of New York, (Spring 2007)
- Gurpreet Kaur, City College of City University of New York, (Spring 2007)
- Bhaveshkumar Patel, City College of City University of New York, (Spring 2007)
- Khurram Malik, City College of City University of New York, (Spring 2007)
- Sandy Sakani, City College of City University of New York, (fall 2006)
- Prasanna Uday Patil, City College of City University of New York, (fall 2006)
- mattupalli susheela, City College of City University of New York, (fall 2006)
- Varun Kumar Yadav Nalla, City College of City University of New York, (fall 2006)
- Qasim Ali Mir, City College of City University of New York, (fall 2006)
- Babji Reddy, City College of City University of New York, (fall 2006)

- Mehul shah, City College of City University of New York, (summer 2006)
- Pramod Yadav, City College of City University of New York, (spring 2006)

◇ **Supervisor of graduate research (independent study) of:**

- Anthony Christopher Lanzisera (NYU, Fall 2020)
- Jeewon Ha (NYU, Fall 2020)
- Siddesh Ramesh (NYU, Fall 2020)
- Chia-hao Wu (NYU, Summer 2020)
- Dhara Mungra (NYU, Spring 2020)
- Tulsi Jain (NYU, Summer 2019)
- Anshu Tomar (NYU, Spring 2019)
- Aviral Khatta (NYU, Fall 2018)
- Arnav Kansal (NYU, Fall 2018)
- Ramani Kothadia (NYU, Fall 2017)
- Yew Wong (NYU, Spring 2017)

◇ **Supervisor of undergraduate research of:**

- Xin Xiang (NYU, Fall 2021)
- Nate Smith (NYU, Fall 2021)
- Andy Huang (NYU, Fall 2020)
- Isabella Tochtermann (NYU, summer 2019)
- Neal K Moorthy (NYU, Fall 2017)
- Shiyang Wang (NYU, Spring 2017)
- Reese Hyde (NYU, Summer 2017)
- Sanjna Verma (NYU, Summer and Fall 2015)
- Katelyn S Mulgrew (NYU, Summer and Fall 2015)
- Daniel Cohen (NYU, Spring, Summer, and Fall 2015)
- Louisa Lu (NYU, Summer and Fall 2014, and Spring 2015)
- Abhi Kumar (CUNY, Spring 2009)
- Lina Cordero (CUNY, Fall 2007)
- Heba Gabre(CUNY, Fall 2007)
- Stephany Soria (CUNY, Fall 2007)
- Jumie Yuventi(CUNY, Fall 2007)
- Jerry Backer (CUNY, Spring 2007, Fall 2007)
- Elbert Tsang (CUNY, Fall 2006)
- Rajai Gooden (CUNY, Spring 06)
- Mamadou Lame (CUNY, Fall 05)
- Juan P. Monzon (CUNY, Spring 05-Fall 05)
- Ruhul Amin (CUNY, Spring 05)
- Mahfuzur Rahman (CUNY, Spring 05)

RESEARCH
SUMMARY

◇ **AI Support for Architecture(2016-present)**

Can we make use of AI techniques to make the hardware more efficient. More specifically, can the execution of program A on a hardware make that hardware more efficient in executing a different program B?

- ◇ **Graphics Processing Units (GPUs)**(2012-present)
In this research project we study the opportunities for using GPUs in many non-scientific applications. We also strive to overcome many challenges ranging from bandwidth bottleneck and power-dissipation, to exposing different types of parallelism.
- ◇ **Memory Design and Management for Heterogeneous Computing** (2016-present)
There are many challenges facing exascale computing. One of them is the heterogeneity involved in computing nodes, memory technology, interconnect, etc. Memory system is one of the crucial parts of the whole system. What technology shall we use for memory modules? And why? How to manage them? What is the role of the compiler, operating system, and the architecture?
- ◇ **Cache Hierarchy Extremely Scalable and Smart** (2006-2015)
In this project, we design the cache hierarchy for many-core chips (i.e. more than 100 on-chip cores). The desired hierarchy must allow cores to exchange information in a fast way, even during very short cycle time, while decreasing off-chip bandwidth requirement. This requires rethinking of many aspects of cache design such as the interaction among caches and the bandwidth wall, the need for a global replacement and placement policy, and the hierarchical decision making in cache hierarchy.
- ◇ **Dynamic Trusted Platform Modules**(2010-2011)
Trusted Platform Module (TPM) is a module placed on-board to provide load-time authentication of general purpose computing systems. This project has three goals. The first is to extend TPM to provide *trusted* execution of programs, and not be limited to load-time or boot-time. The second is to solve the scalability issues of TPMs. The third goal is to embed TPMs into multicore and manycore chips for tighter interaction with the different cores.
- ◇ **Chip Multithreading**(2004-2010)
As the number of transistors has already exceeded billion transistors per chip, the inclusion of several processing elements per chip is now feasible. Making the best use of this chip-multiprocessing capability is a challenging problem. To achieve that goal, we study problems such as: dynamic balancing between high throughput and high single-thread performance, fast and dynamic pattern detection of different application behaviors, and the dynamic reconfiguration of several processing elements on-chip.
- ◇ **Reconfigurable and Parallel Architecture:**(2003-2004)
In this project, we study the integration of both general-purpose processors as well as reconfigurable engines in order to allow the machine to dynamically reconfigure itself based on run-time information to achieve the performance of an ASIC with the flexibility of a general purpose machine. Moreover, the project involves the design of language extension to the C-programming language for shared memory architectures. The language is called Unified Parallel C (UPC).
- ◇ **Hierarchical Multithreading** (1999-2003)
This is a new execution model I proposed in my Ph.D thesis that makes use of speculative multithreading in a novel way. It is used to parallelize a sequential program using a combined hardware and software approaches.