Snehasish Kumar

mail@snehasish.net
snehasish.net
github.com/snehasish

Interests

- Profile guided optimizations
- Hardware software co-design for specialization
- Micro-architecture research with a focus on the memory hierarchy

Professional Experience

08/17 – Staff Software Engineer : Core, Google
Profile guided heap layout optimizations
Post-link optimizations (published ASPLOS'23)
Workload analysis using performance counters and dynamic tracing tools

06/13 – 12/13 Research Intern : IBM, T.J. Watson Research Centre
Compiler based analyses of loops for hardware specialization (published ICS'16)

01/11 – 11/16 Research Assistant : Computer Architecture Research Group, Simon Fraser University
Design of LLVM IR abstractions for hardware accelerators (published MICRO'16)
Implementation of low overhead, dynamic profiling for characterization (published IISWC'16)
Interfacing LLVM IR to FPGA (Verilog) code generators for irregular programs (published HPCA'17)
Design and evaluation of coherence protocols for hardware accelerators (published ISCA'15)
Design and evaluation of variable granularity caching mechanisms (published MICRO'12, ISCA'13)

Academic Qualifications

05/13 – 11/16 **PhD in Computing Science**, *Simon Fraser University*, British Columbia, Canada, *4.0/4.0*

01/11 – 04/13 **MSc in Computing Science**, *Simon Fraser University*, British Columbia, Canada, *3.8/4.0*

08/06 – 04/10 **B. Tech in Computer Engineering**, *BPUT*, Orissa, India, *8.3/10.0*

Awards

08/16 President's PhD Scholarship, Simon Fraser University

'16, '14, '12 Graduate Fellowship, Simon Fraser University

01/14 Special Graduate Entrance Scholarship, Simon Fraser University

Publications

2023 – **Propeller: A Profile Guided, Relinking Optimizer for Warehouse-Scale Applications,**
Han Shen, Krzysztof Pszeniczny, Rahman Lavaee, Snehasish Kumar, Sriraman Tallam, and Xinliang (David) Li,
28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS '23.
Acceptance Rate ≈ X%.

2017 – **Needle: Leveraging program analysis to extract accelerators from whole programs,**
Snehasish Kumar, Nick Sumner, Vijayalakshmi Srinivasan, Steve Margerm, and Arrvindh Shriraman,
23rd ACM International Conference on High Performance Computer Architecture, HPCA '17.
Acceptance Rate ≈ 22%.

2016 – **ChainSaw: Creating Von-Neumann Accelerators with Fused Instruction Chains,**
Amirali Sharifian, Snehasish Kumar, Apala Guha, and Arrvindh Shriraman,
49th Annual IEEE/ACM International Symposium on Microarchitecture, MICRO '16.
Acceptance Rate ≈ 22%.

– **SPEC-AX: Extracting Accelerator Benchmarks from Microprocessor Benchmarks,**
Snehasish Kumar, Nick Sumner, and Arrvindh Shriraman,
2016 IEEE International Symposium on Workload Characterization, IISWC '16.
Acceptance Rate ≈ 30%.

– **Peruse and Profit: Estimating the Accelerability of Loops,**
Snehasish Kumar, Vijayalakshmi Srinivasan, Amirali Sharifian, Nick Sumner, and Arrvindh Shriraman,
30th ACM International Conference on Supercomputing, ICS '16.
Acceptance Rate ≈ 24%.

2015 – **Fusion: Design Tradeoffs in Coherent Cache Hierarchies for Accelerators,**
Snehasish Kumar, Arrvindh Shriraman, and Naveen Vedula,
42nd Annual International Symposium on Computer Architecture, ISCA '15.
Acceptance Rate ≈ 19%.

– **DASX: Hardware Accelerator for Software Data Structures,**
Snehasish Kumar, Naveen Vedula, Arrvindh Shriraman, and Vijayalakshmi Srinivasan,
29th ACM International Conference on Supercomputing, ICS '15.
Acceptance Rate ≈ 25%.

2013 – **Protozoa: Adaptive Granularity Cache Coherence,**
Hongzhou Zhao, Arrvindh Shriraman, Snehasish Kumar, and Sandhya Dwarkadas,
40th Annual International Symposium on Computer Architecture, ISCA '13.
Acceptance Rate ≈ 19%.

2012 – **Amoeba-Cache: Adaptive Blocks for Eliminating Waste in the Memory Hierarchy,**
Snehasish Kumar, Hongzhou Zhao, Arrvindh Shriraman, E. Matthews, S. Dwarkadas, and L. Shannon,
45th Annual IEEE/ACM International Symposium on Microarchitecture, MICRO '12.
Acceptance Rate ≈ 18%.

Workshops

03/16 GCASR'16 – Statistical program analysis assisted cost-effective sampling in large scale scientific simulations
06/15 SFU-ZU workshop on Big Data – Data Structure Accelerators
12/13, 08/14 WoNDP'13, PACT'14 – SQRL: Hardware Accelerator for Collecting Software Data Structures

Invited Talks

06/16 IBM Research – Needle [HPCA '17]
01/16 SRC India Design Review – Caches [MICRO '12, ISCA '15]
01/16 Intel Bangalore – Fusion [ISCA '15]

Service

CGO '20, '23 Technical Program Committee
PPoPP '18, '19 Artifact Evaluation Committee
PLDI '17 Artifact Evaluation Committee
HiPC '20 Technical Program Committee

Projects

01/15 Networks: Parallel implementation of Kou, Markowsky and Berman (1981) algorithm

04/14 Natural Language Processing: Optimizing the Bitpar CKY parser

12/11 Computational Geometry: Interactive demo for the Linear Cell Complex (CGAL)

04/11 Machine Learning: Non-Negative Matrix Factorisation for large datasets