
Ph.D. @CSE, IITH

CSE IITH-PHD ADMISSIONS OPEN 2025

భారతీయ సాంకేతిక విజ్ఞాన సంస్థ హైదరాబాద్
भारतीय प्रौद्योगिकी संस्थान हैदराबाद
Indian Institute of Technology Hyderabad

CSE DEPARTMENT HIGHLIGHTS

Broad Research Areas

01 Theory
- Algorithms
- Computational Complexity
- Graph Theory
- Combinatorics
- Formal Methods
- Quantum Computing
- Post-quantum cryptography

02 Systems
- Computer Networks
- Compilers
- Architecture
- Distributed Systems
- Blockchains
- Cyber Security

03 AI/ML & Data Science
- Computer Vision
- NLP (Natural Language Processing)
- Social Media Analytics
- Theoretical AI/ML
- Applications

PhD Fellowships

- Google Ph.D. Fellowship
- TCS Ph.D. Fellowship
- Intel Ph.D. Fellowship
- Reliance Foundation Fellowship
- PMRF Fellowship
- Microsoft Research India PhD Award
- Qualcomm Innovation Fellowship

Alumni in PostDoc Positions

Publication Venues

Student Recognitions

- IDRBT Doctoral Colloquium
- Indo-Canadian Shastri Student Research Fellowship
- S N Bose Fellowship
- Honda YES Fellowship
- Viterbi Fellowship
- Google AI Residency
- Facebook AI Residency
- Fulbright-Nehru doctoral research fellowship

Eligibility

BE (or equivalent) with a valid GATE score, or students from a CFTI with a CGPA of 8.0 and above, or MTech degree holders can apply.

PhD Programme
- MTech
- B.Tech (or equivalent) with a valid GATE score
- B.Tech(or equivalent) from CFTI with CGPA 8.00 and above

How to Apply

How to Apply and Prepare?

- Applications must be submitted online at: https://iith.ac.in/phdadmissions/
- CSE PhD information at: https://cse.iith.ac.in/admissions/phd.html
- For PhD admission related queries: Click here
- How to Prepare for Ph.D. Interviews? Click here
- Syllabus and Sample questions: Click here
- Current PhD Students: Click here

Selection process: May 2025
(exact dates will be informed via email)

Application start date : 12 March 2025
Application end date : 10 April 2025

CSE department brochure:
https://cse.iith.ac.in/aboutUs/Brochure.html

Table of Contents

Theoretical Computer Science @IITH

Cryptography

- How to efficiently authenticate a vehicle without revealing private information in fast moving traffic? Lightweight cryptography!

Cryptography

- A quantum computer can break most of the encryption schemes of today. What are the options for crypto, post-quantum?
  Quantum crypto, Lattice crypto, etc.

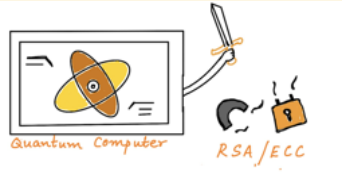

- Hardware/Software aspects of quantum cryptography.
- Privacy preserving mechanisms over blockchains.

Combinatorics

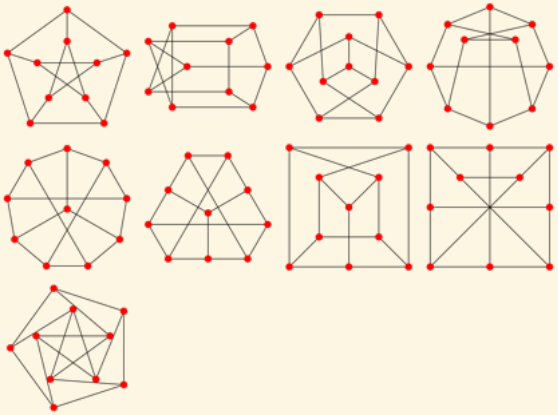

Are they the same graph?

Graph Isomorphism Problem

(Image Courtesy: Wolfram Mathworld)

Combinatorics

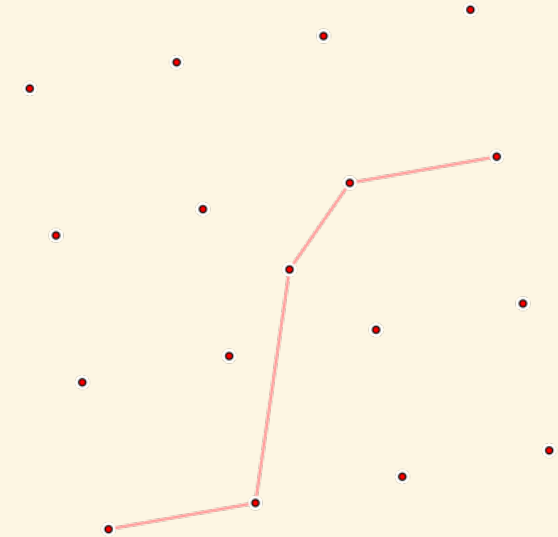

(Image Courtesy: David Eppstein (from Wikipedia))

Erdős Szekeres Problem:

What is the longest increasing/decreasing sequence here?

Combinatorics

(Image Courtesy: BIG MOUTH for Quanta Magazine)

Second Neighborhood Problem:

Posed by Paul Seymour (1990):
In a social network described by such a graph, is there always someone who has at least as many friends-of-friends as friends?

Algorithms

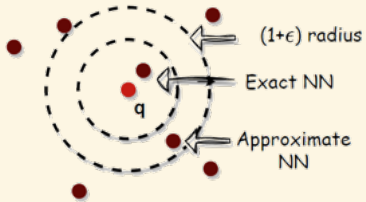

Nearest Neighbors (NN) Search

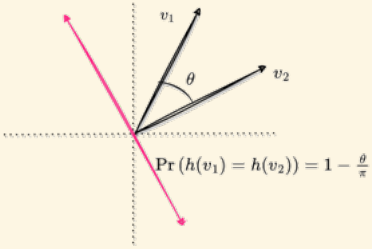

Nearest Neighbors Search via Hashing

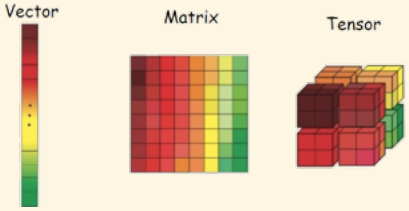

Image courtesy: kdnuggets

- Locality Sensitive Hashing (LSH) [Indyk, Motwani '98] suggest approximate nearest neighbour search algorithm for vectors.
- Major open problem is to propose (approximate) nearest neighbour search algorithms for tensors!

Algorithms

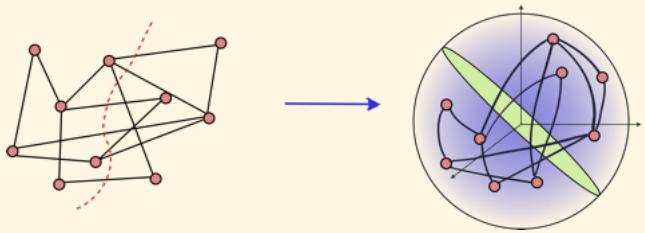

Max-Cut problem:

Partition a set of people into two parts such that interaction across parts is maximized.

[Goemans-Williamson ‘95]:
Embedding the graph into a sphere, and cutting the sphere into two halves to find the partition gives a good approximation.

Major open problem to find a better solution!

Complexity Theory

Travelling Salesman Problem:
Given: cost of travelling between every pair of cities, and a cost c.
Is there a tour with cost $\leq c$ that visits every vertex exactly once and ends in the starting vertex?

Image Courtesies: Media from chegg.com, and linkedin.com

Complexity Theory

Zero Knowledge Proofs:
Can you prove that you have found Waldo without revealing where he is?

Image Courtesies: firstpost.com, and pngitem.com

Theoretical Computer Science Faculty

Maria Francis

Cryptography, Computational Algebra

Rogers Mathew

Combinatorics

M. V. Panduranga Rao

Quantum Computing

Nitin Saurabh

Computational Complexity, Algorithms

Rakesh Venkat

Algorithms

Subrahmanyam Kalyanasundaram

Computational Complexity

Aravind N.R

Graph Theory, Algorithms, Combinatorics

Rameshwar Pratap

Algorithms, Machine Learning

Karteek Sreenivasaiah

Computational Complexity, Algorithms

Networks and Systems Group
Dept. of Computer Science and Engineering

భారతీయ సాంకేతిక విజ్ఞాన సంస్థ హైదరాబాద్
भारतीय प्रौद्योगिकी संस्थान हैदराबाद
Indian Institute of Technology Hyderabad

Faculty

- Saurabh Kumar

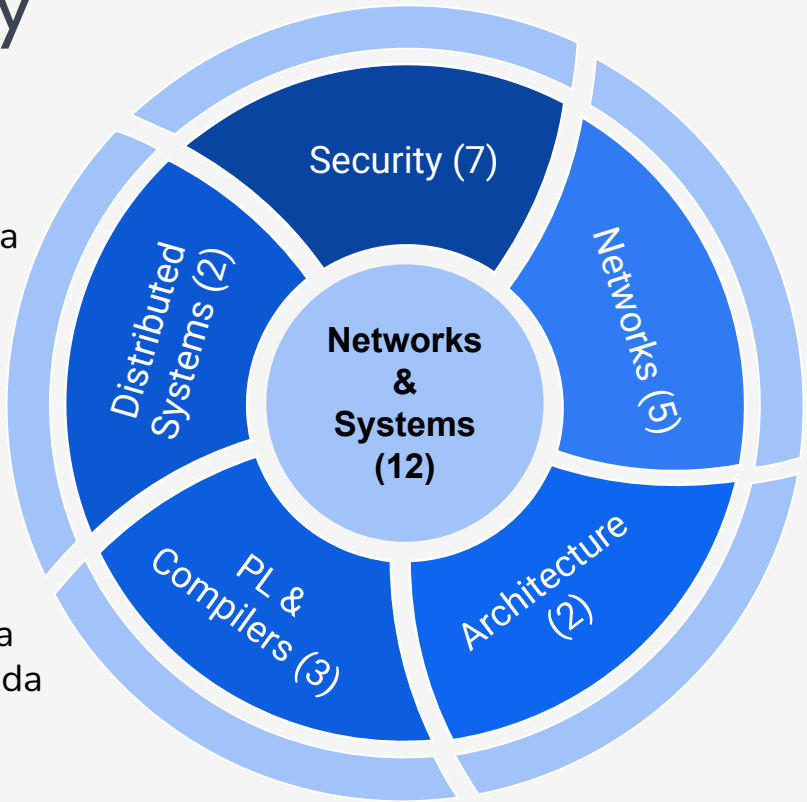

- Antony Franklin
- Bheemarjuna Reddy Tamma
- C. Siva Ram Murthy
- Kotaro Kataoka
- Praveen Tammana

- Kotaro Kataoka
- Sathya Peri

- Rajesh Kedia
- Shirshendu Das

- Ashish Mishra
- Jyothi Vedurada
- Ramakrishna Upadrasta

Research Expertise: Networks & Security

01	**Mobile Wireless Networks (5G and Beyond Networks)**
02	Software-defined Networking & Network Functions Virtualization
03	Networked systems for AI/ML: Programmable Data Planes, Data Center Networking
04	V2X & Mobile Edge for Autonomous Navigation
05	Cyber Security, Wireless Security & AI





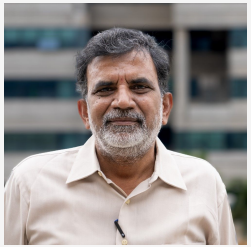


Research Expertise: PL/Compilers

01 **Program Analysis and Compilers using Machine Learning**

02 Polyhedral Compilation

03 Compiler Optimization Techniques for CPUs & GPUs

04 High-performance GPU Algorithms for Scalable Graph Analysis

05 Data Race Checking & Parallel Code Compliance Standards

06 Tree Automata for Synthesis and Verification over Refinement Typed Space

07 ML Modulo Repair for scalable synthesis.

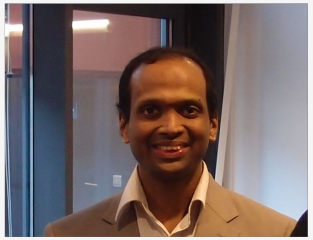

Research Expertise: Computer Architecture

01 **Thermal Management for 3D Architectures**

02 Memory Security

03 Cache, Interconnects & Memory Access Policies

04 Resource Sharing in Heterogeneous Architectures

05 Non-volatile Memories

Research Expertise: Distributed Systems

01 **Blockchains & its applications to Security and Smart Contracts**

02 Large-Scale Graph Analytics

03 Efficient Consensus Protocols

04 Distributed and Federated Learning

05 Lock-Free Programming

Courses Offered: Electives

1. Networked Wireless Systems
2. Internet of Things
3. Topics in Networks (/w Industry)
4. Topics in Wireless Networks
5. Software Defined Networks
6. Network Engineering (hands-on)
7. **Cyber Security and AI (/w Industry and AI dept.)**
8. Concurrency Control in Transactional Systems
9. Parallel and Concurrent Programming
10. Distributed Computing
11. Parallel Programming for Practitioners
12. Compiler Optimizations
13. Advanced Compiler Design
14. Advanced Compiler Optimizations
15. Topics in Compiler Optimizations 🤝
16. Introduction to Compiler Engineering (/w Industry)
17. Advanced Compiler Engineering
18. Compilers for Machine Learning
19. Introduction to Program Analysis and Compiler Optimization
20. **Program Synthesis**
21. Advanced Computer Architecture 🤝
22. **Hardware Architecture for Deep Learning** 🤝

Research Highlights

Networked systems research

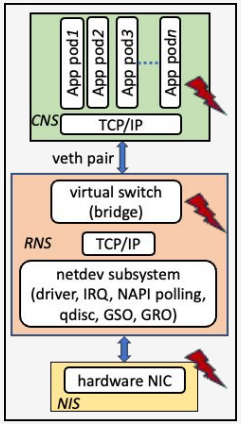

Host network stack (eBPF, smartNICs)

1. Monitoring and detecting performance anomalies
Challenge: Minimizing per-packet overheads at high-speed (> 100Gbps)

2. Static analysis of eBPF bytecodes for debugging
Challenge: Extracting network context to bytecode

Collaborators: IITH Compilers team, IBM Res, CMU

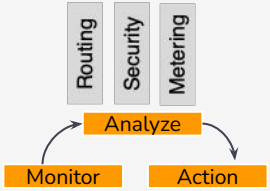

- AI network cluster
- Data center network
- Private WAN

1. Traffic Engineering (P4, SDN):
- Congestion control, multipath traffic split, In-network telemetry
Challenge: Restrictions on per-packet operations, state, latency

2. Securing in-network systems (P4, SDN):
- DoS, DDoS, Evasion, Poisoning, Bad QoS
Challenge: Tradeoff between performance and security operations

Collaborators: IIIT Delhi, Princeton, Univ. of Waterloo, Marvell

- Edge cloud offload
- Cloud-native networking

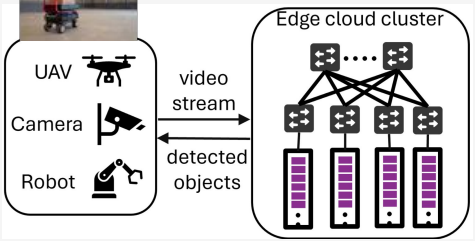

1. Load-aware in-network LB (service mesh)
Challenge: Visibility into servers' load and nw path status

2. GPU-centric packet processing for IoT-ML inference –GPUNetIO, GPUDirectRDMA
Challenge: Reliable transport over lossy networks

Collaborators: NYU, IITH, DigiQuanta

Department of Science & Technology, Government of India
सत्यमेव जयते

Edge Cloud for Autonomous Navigation Applications

TiHAN IIT HYDERABAD
भारतीय सांकेतिक विज्ञान पंथ हैदराबाद
भारतीय प्रौद्योगिकी संस्थान हैदराबाद
Indian Institute of Technology Hyderabad

Goal: Bring compute and storage closer to the data source

TiHAN testbed for Research & Technology development of Autonomous Navigation and Data Acquisition Systems

TiE Architecture

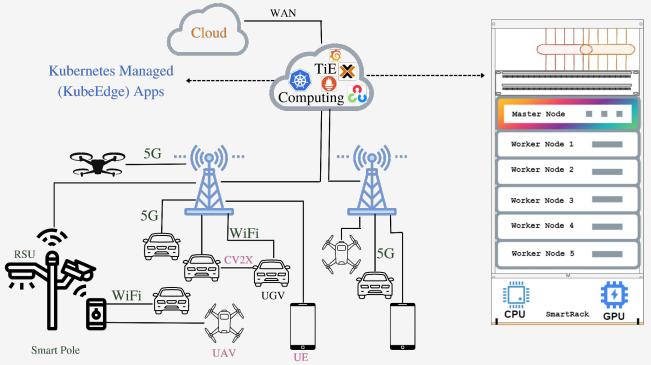

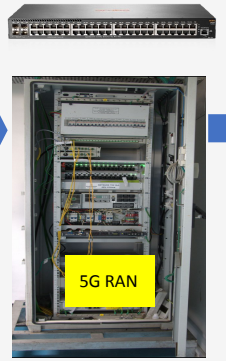

"Experience: Implementation of Edge-Cloud for Autonomous Navigation Applications" IEEE/ACM COMSNETS, Bangalore, India.
Link1, Link 2

OAI based 4G/5G Virtualized Base Station Testbed

- Joint allocation of radio and system resources
- Traffic-aware compute resource tuning for energy efficiency
- A flexible split based 5G RAN to Minimize energy consumption and handovers

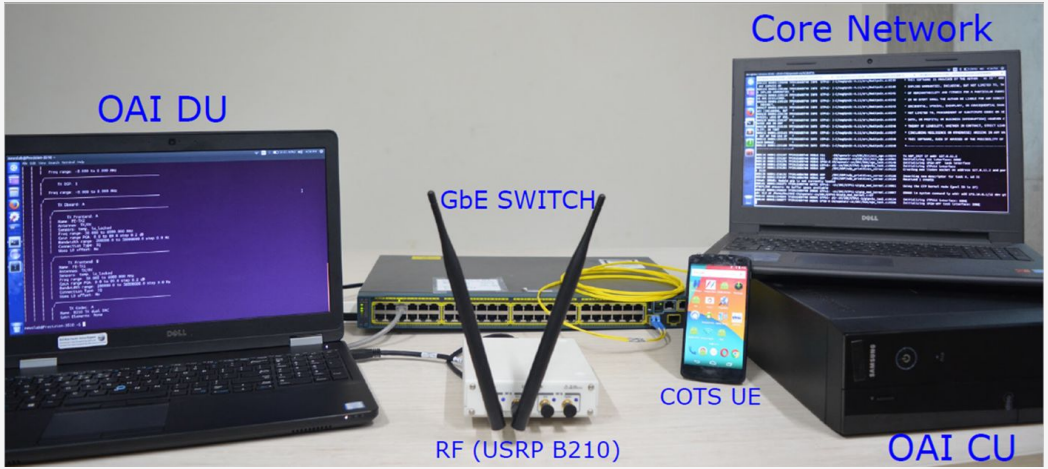

[IEEE Netsoft] [IEEE TNSM] [IEEE NOMS]

Building IITH 5G Core

- Funded by “Indigenous 5G Testbed” Project, DoT, GoI
- Supports Network Slicing and Orchestration (OSM)
- Integration with OAI RAN & MEC platforms
- Contributing to open source platforms like OpenAirInterface (OAI)

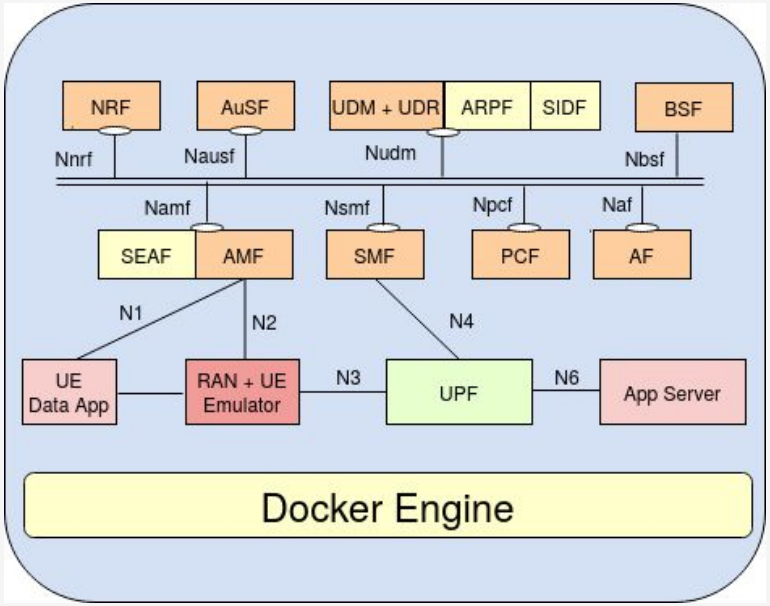

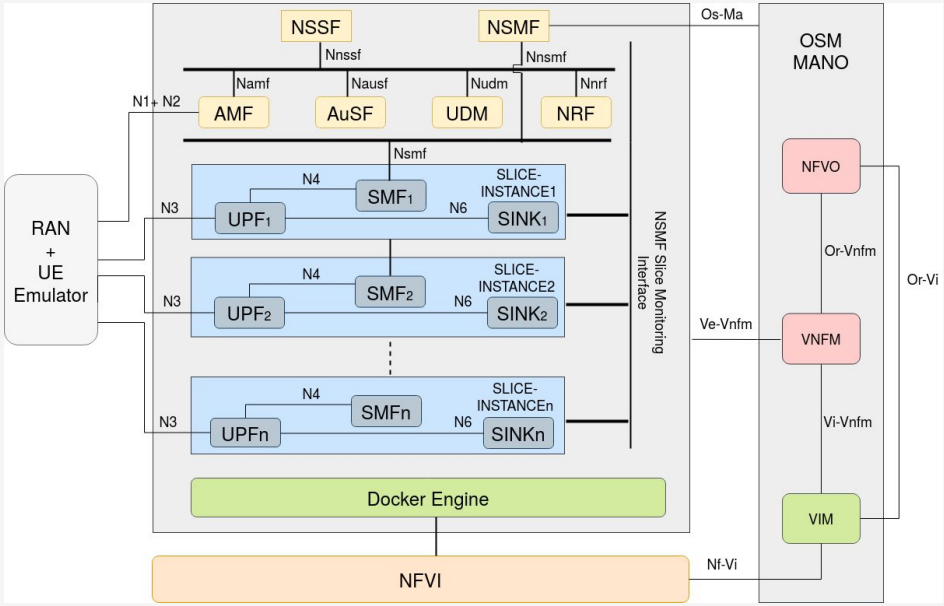

Zoom-in on some of the Developed Frameworks

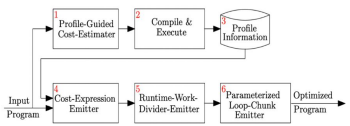

Deep Chunking: Chunking Loops with non-Uniform Workloads

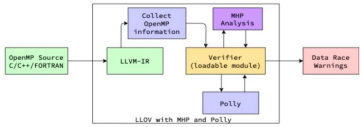

LLOV: Static Data-Race Checker for OpenMP Programs

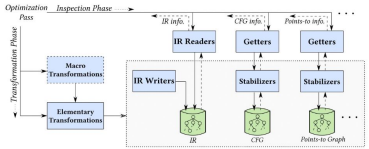

Homeostasis: Self-Stabilizing Compiler Infrastructure

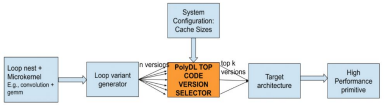

PolyDL: Polyhedral Optimizations for Creation of High-performance DL Primitives.

Papers Published

- DisGCo : A Compiler for Distributed Graph Analytics. R S Anchu and V K Nandivada, *in the ACM Transactions on Architecture and Code Optimization (TACO)*, 2020. Presented at HiPEAC 2021.
- Chunking Loops with non-Uniform Workloads. I K Prabhu and V K Nandivada, *in the Proceedings of the International conference on Supercomputing (ICS)*, 2020.
- PolyDL: Polyhedral Optimizations for Creation of High-performance DL Primitives. Sanket Tavarageri, Alexander Heinecke, Sasikanth Avancha, Bharat Kaul, Gagandeep Goyal, and Ramakrishna Upadrasta. 2021. *ACM Trans. Archit. Code Optim.* (TACO), 2021. Presented at HiPEAC 2021.
- LLOV: A Fast Static Data-Race Checker for OpenMP Programs. Utpal Bora, Santanu Das, Pankaj Kukreja Saurabh Joshi, Ramakrishna Upadrasta, and Sanjay Rajopadhye. 2020. *ACM Trans. Archit. Code Optim. (TACO),* 2020. Presented at HiPEAC 2021.
- OpenMP aware MHP Analysis for Improved Static Data-Race Detection Utpal Bora, Shraiysh Vaishay, Saurabh Joshi, Ramakrishna Upadrasta. The Seventh Annual Workshop on the LLVM Compiler Infrastructure in HPC Workshop held in conjunction with SC21 - November 14, 2021 - St. Louis, USA

Software Produced (or *planned* to be Released)

- **IMOP:** IIT Madras OpenMP Compiler Framework (http://www.cse.iitm.ac.in/~amannoug/imop/)
- **DiscGo:** A compiler to compile Green-Marl programs to efficient MPI Programs (https://github.com/anchur/DisGCo)
- **LLOV (2.0):** LLVM OpenMP Verifier https://compilers.cse.iith.ac.in/projects/llov/
- **PolyDL/PolyAI:** Polyhedral Optimizations for Creation of High Performance DL primitives https://compilers.cse.iith.ac.in/projects/polydl/
- **BullsEye:** Scalable and Accurate Approximation Framework for Cache Miss Calculation https://compilers.cse.iith.ac.in/projects/bullseye/

Tutorials of Softwares

- **IMOP**
  - CGO 2020
  - CGO 2021, CGO 2022 (online)
  - (upcoming) CGO 2023
- **LLOV:** Fourth LLVM performance workshop (Co-located with CGO-2020)
- **PolyDL/PolyAI:** Compilers for Machine Learning (C4ML) workshop 2021 (co-located with CGO-2021)
- **POSET-RL:** Sixth LLVM performance workshop (Co-located with CGO-2022) (online)

MEC Platform for Deploying 5G Applications

- Integrated /w IITH 5GC
- Support for Radio Network & Location Information Services
- Edge video caching for DASH-based video streaming

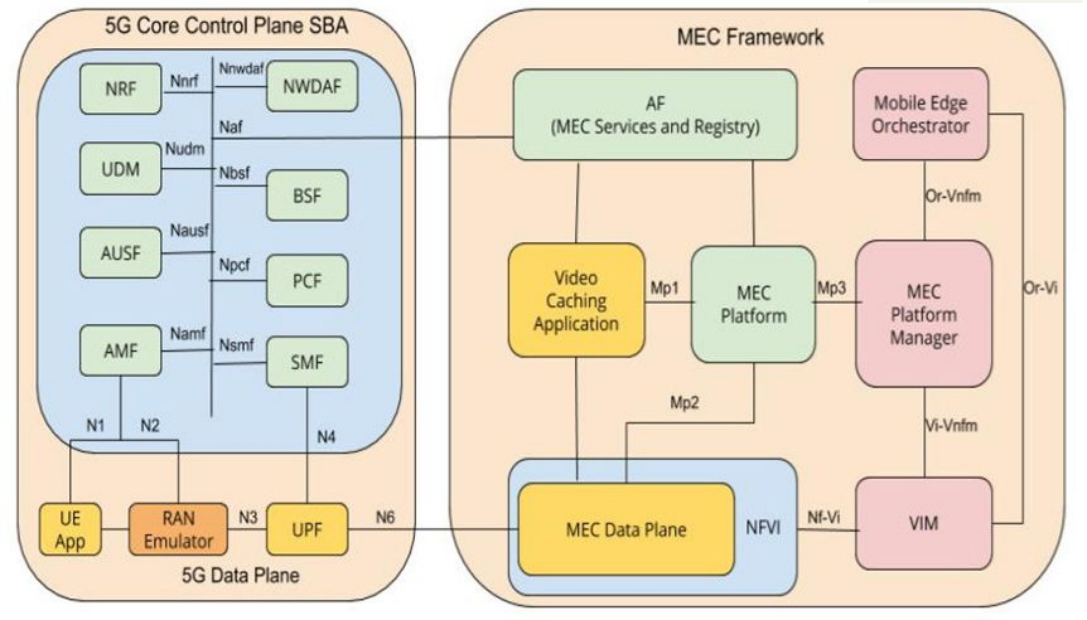

[IEEE Netsoft] [IEEE/ACM COMSNETS]

Continual Learning for Network Intrusion Detection

- Severe class imbalance (CI) in datasets like CICIDS-18, CTU-13 and ANOSHIFT
- Scalability issues in using CL for realtime NIDS
- Extended Class based Reservoir Sampling (CBRS) for tackling severe CI
- Perturbation assistance for parameter approximation (PAPA) for scalability

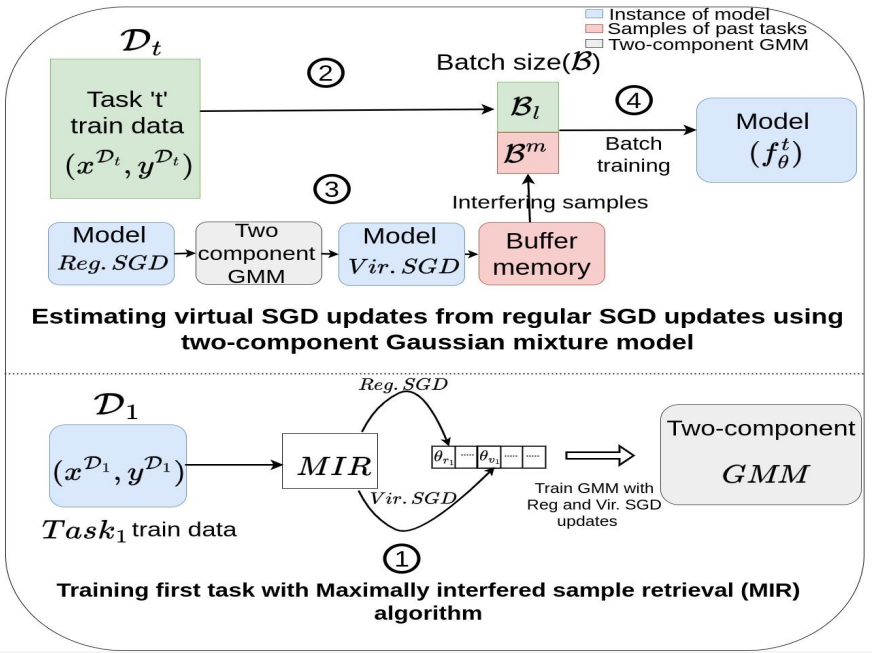

[NeurIPS 2023]

Internet Architecture

- Digital Twin of SDN based Networks by Data Representation
- Topology Poisoning Attacks and Prevention in Hybrid SDN (IEEE TNSM 2022)
- Slow HTTP DoS Attack Detection using Autoencoders through Unsupervised Learning (ACM AINTEC 2021)

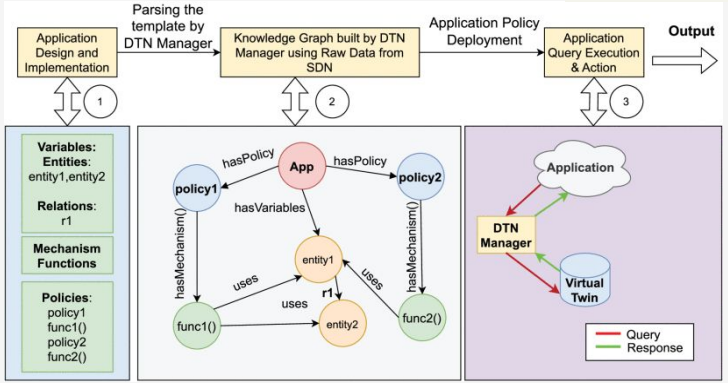

Workflow of Constructing a Digital Twin Network

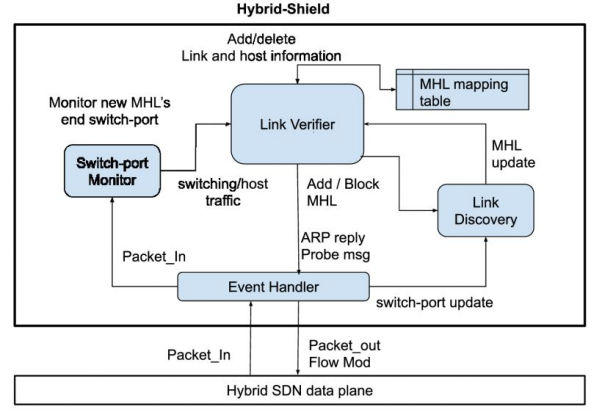

Workflow of Prevention Mechanism of Topology Poisoning Attacks

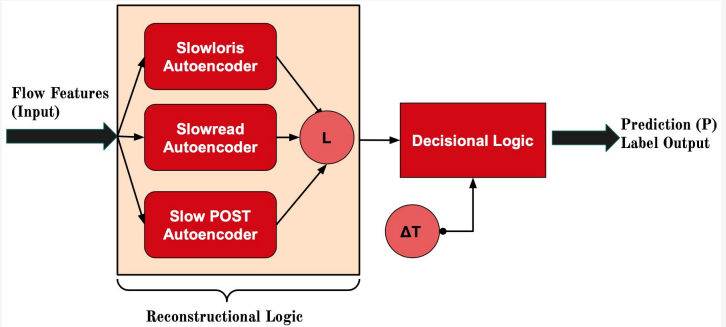

Slow HTTP DoS Attack Detection

Blockchain

- DAG based Distributed Ledger for IoT Data Integrity (IEEE ICOIN 2021)
- Collective Intelligence by AI-Blockchain Interplay (IEEE IV 2023)
- Credential Management (IEEE Access 2022)
- Inter Blockchain Communication (ongoing)
- Exploration of Blockchain Use Cases in Mental Health Care
- Industry Collaborators: DENSO, I'mbesideyou, Chaintope, etc.

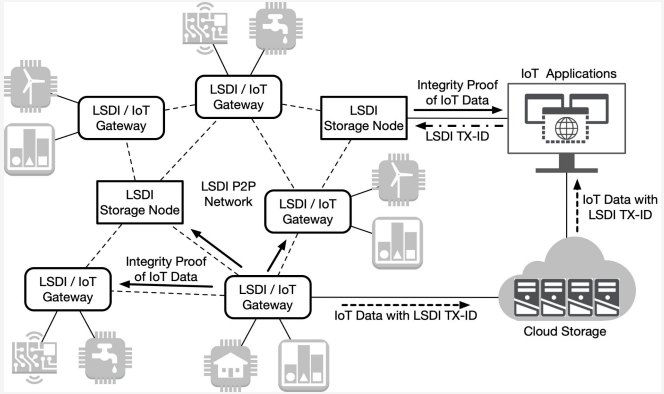

Working of DAG based Distributed Ledger

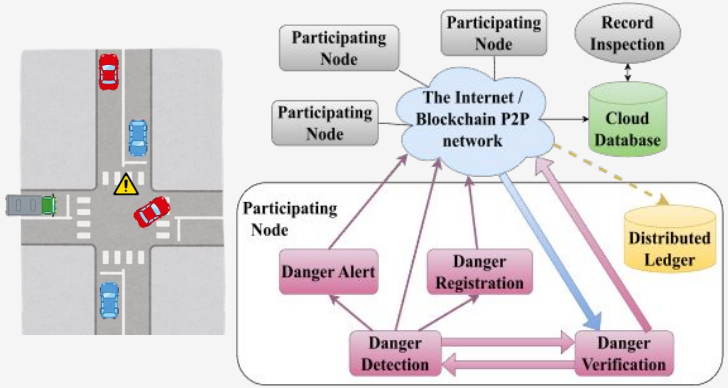

Collective Intelligence by AI-Blockchain Interplay

Parallel & Distributed Systems

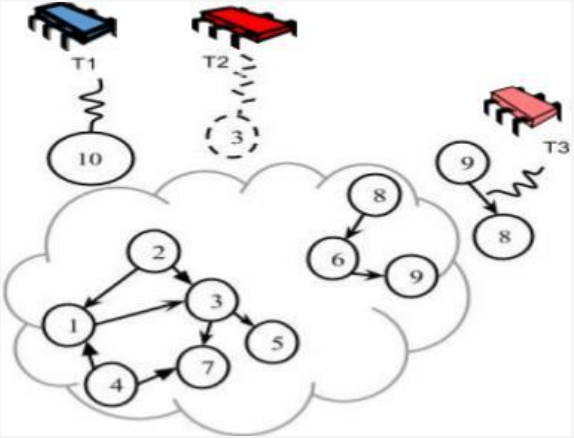

Illustration of a Dynamic Time-Changing Graph

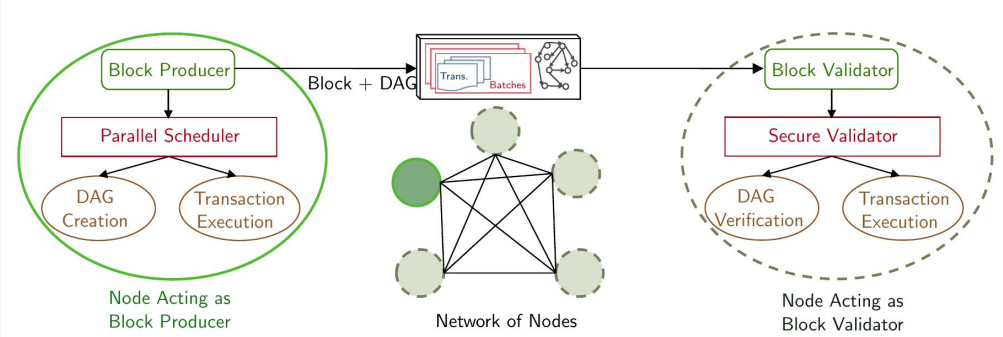

Efficient Smart Contract Execution in Blockchains

Dynamic Distributed Graph Analytics

- Working on developing tools for Dynamic time-changing distributed graphs
- Worked the analytics operations: Betweeness Centrality, Pagerank, BFS, Shortest Paths

Efficient Smart Contract Execution in Blockchains

Developed an Efficient Distributed and Secure framework for the Execution of Smart Contracts in Blockchains.
Demonstrated it Hyperledge Sawtooth

Computer Architecture

Problem: Standard interconnect and memory management policies are not efficient for accelerator-rich systems

Research: Identify resource sharing policies considering accelerator access patterns

Problem: 3D architectures for processor and memory suffer from thermal issues

Research: Propose thermal management policies with low overheads

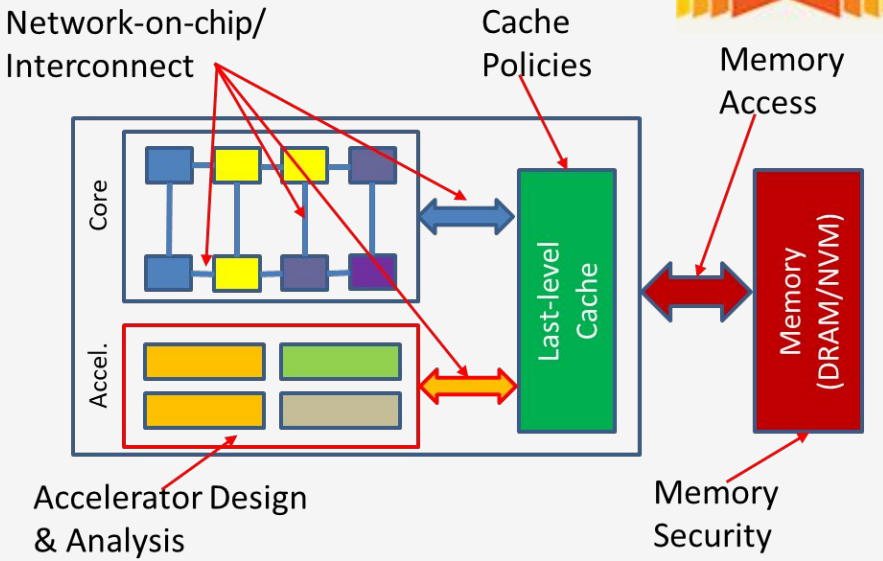

Computer Architecture

Concurrent CNN Accelerators

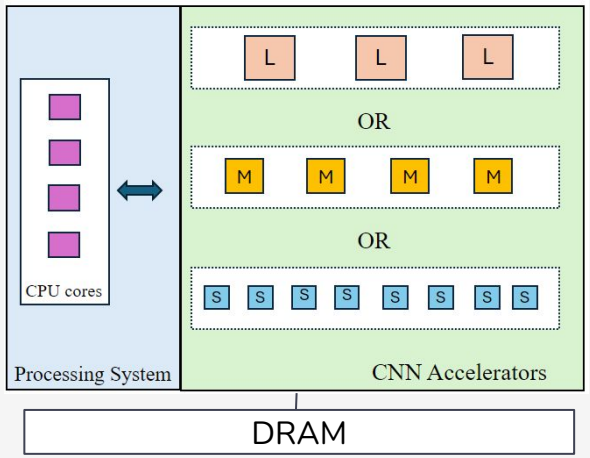

- Characterizing workloads for various hardware configurations [IEEE ESL 2024]
- CNN execution time estimation as per hardware configuration [ACM TECS 2024]

RISC-V Specific Optimizations

- Custom/co-processor instructions as per application requirements
- Register access optimizations - hardware and software

CNN/LLM Optimizations for NVM

- NVM have high latency and energy cost for writes
- Analyze the bit changes and optimize the total number of writes

Chipmultiprocessor and Hardware Security

In modern Chipmultiprocessors (CMP) multiple cores share a common Last Level Cache (LLC)

- Performance enhancement of CMP LLCs: *Replacement Policies, Prefetching, Cache Utilisation, and Coherence Issues.*
- Preventing CMP LLCs from Side-Channel and Covert-Channel attacks without compromising on performance.
- Exploring the possibilities of different Hardware Trojan (HT) present in the Network-on-Chip (NoC) of the modern CMPs.
- Resolving the challenges of using emerging memory technologies to design CMP LLC: *Performance, Endurance and Security.*
- Reducing the refresh overhead and preventing the security attacks on DRAM-based main memories.

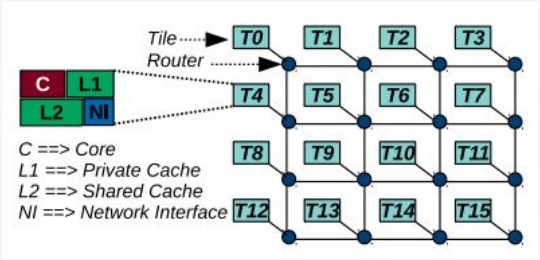

A Chipmultiprocessor (CMP) with tiled architecture.

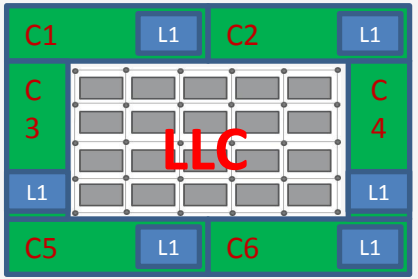

A Chipmultiprocessor (CMP) with non-tiled architecture.

Programming Model/Language, Compiler and Runtime System for Emerging HPC Systems

Zoom-in on some of the Developed Frameworks

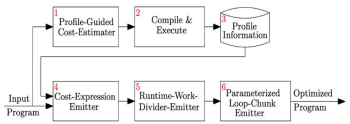

Deep Chunking: Chunking Loops with non-Uniform Workloads

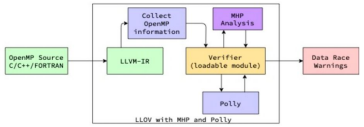

LLOV: Static Data-Race Checker for OpenMP Programs

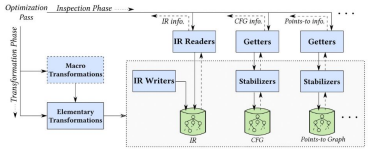

Homeostasis: Self-Stabilizing Compiler Infrastructure

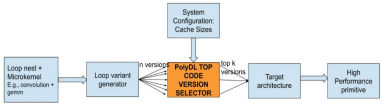

PolyDL: Polyhedral Optimizations for Creation of High-performance DL Primitives.

Papers Published

- DisGCo : A Compiler for Distributed Graph Analytics. R S Anchu and V K Nandivada, *in the ACM Transactions on Architecture and Code Optimization (TACO)*, 2020. Presented at HiPEAC 2021.
- Chunking Loops with non-Uniform Workloads. I K Prabhu and V K Nandivada, *in the Proceedings of the International conference on Supercomputing (ICS)*, 2020.
- PolyDL: Polyhedral Optimizations for Creation of High-performance DL Primitives. Sanket Tavarageri, Alexander Heinecke, Sasikanth Avancha, Bharat Kaul, Gagandeep Goyal, and Ramakrishna Upadrasta. 2021. *ACM Trans. Archit. Code Optim.* (TACO), 2021. Presented at HiPEAC 2021.
- LLOV: A Fast Static Data-Race Checker for OpenMP Programs. Utpal Bora, Santanu Das, Pankaj Kukreja Saurabh Joshi, Ramakrishna Upadrasta, and Sanjay Rajopadhye. 2020. *ACM Trans. Archit. Code Optim. (TACO),* 2020. Presented at HiPEAC 2021.
- OpenMP aware MHP Analysis for Improved Static Data-Race Detection Utpal Bora, Shraiysh Vaishay, Saurabh Joshi, Ramakrishna Upadrasta. The Seventh Annual Workshop on the LLVM Compiler Infrastructure in HPC Workshop held in conjunction with SC21 - November 14, 2021 - St. Louis, USA

Software Produced (or *planned* to be Released)

- **IMOP:** IIT Madras OpenMP Compiler Framework (http://www.cse.iitm.ac.in/~amannoug/imop/)
- **DiscGo:** A compiler to compile Green-Marl programs to efficient MPI Programs (https://github.com/anchur/DisGCo)
- **LLOV (2.0):** LLVM OpenMP Verifier https://compilers.cse.iith.ac.in/projects/llov/
- **PolyDL/PolyAI:** Polyhedral Optimizations for Creation of High Performance DL primitives https://compilers.cse.iith.ac.in/projects/polydl/
- **BullsEye:** Scalable and Accurate Approximation Framework for Cache Miss Calculation https://compilers.cse.iith.ac.in/projects/bullseye/

Tutorials of Softwares

- **IMOP**
  - CGO 2020
  - CGO 2021, CGO 2022 (online)
  - (upcoming) CGO 2023
- **LLOV:** Fourth LLVM performance workshop (Co-located with CGO-2020)
- **PolyDL/PolyAI:** Compilers for Machine Learning (C4ML) workshop 2021 (co-located with CGO-2021)
- **POSET-RL:** Sixth LLVM performance workshop (Co-located with CGO-2022) (online)

ML4Code: Machine Learning for Compiler Optimizations and Program Analysis

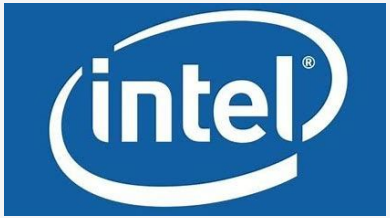

Machine Learning for Code: Embeddings and Applications

IR2Vec - Program Embeddings

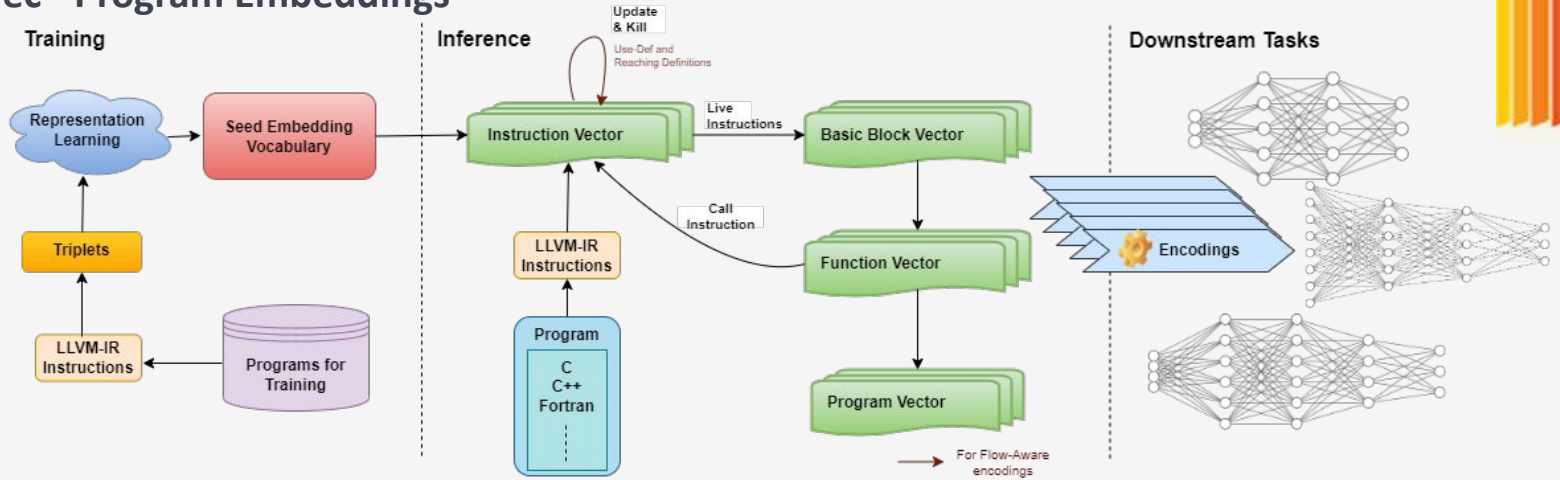

Applications

Heterog enous Device Mapping [TACO' 20, APNET' 22]		**Thread Coarsening** **[TACO' 20]**
Register Allocation [CC' 23]		**Compiler Phase Ordering** **[ISPASS' 21]**
Power and Energy Optimizations [Ongoing]		**Binary and Source code Similarity** **[Ongoing, Under Submission]**

Compiler Optimizations in CPU-GPU Heterogeneous Systems

Before Optimization

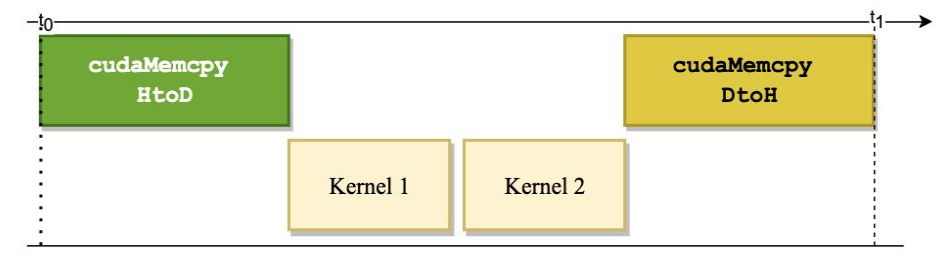

After *StreamAlloc* Optimization

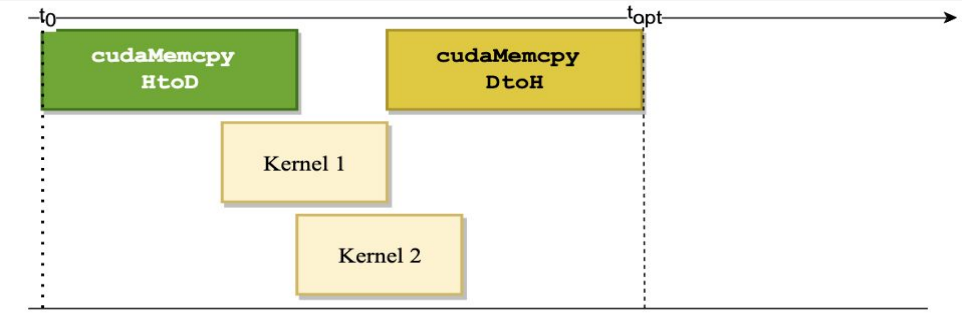

- Hetero-sync Motion Optimization
- Compiler Optimizations to percolate blocking/Sync statements to a later program point to enable overlapping computations between CPU and GPU increasing parallelism.
- Automatic CUDA Stream Allocation
- Compiler Optimization to convert Sync to Async calls in CPU-GPU Heterogeneous programs
- Determines the stream identifiers for each corresponding asynchronous call

Billion-Scale ANN on GPU

- A high-throughput billion-point scale Approximate Nearest Neighbour search on CPU-GPU Heterogeneous system
- On a single A100 GPU
- Graph index remains on host memory but optimized ac

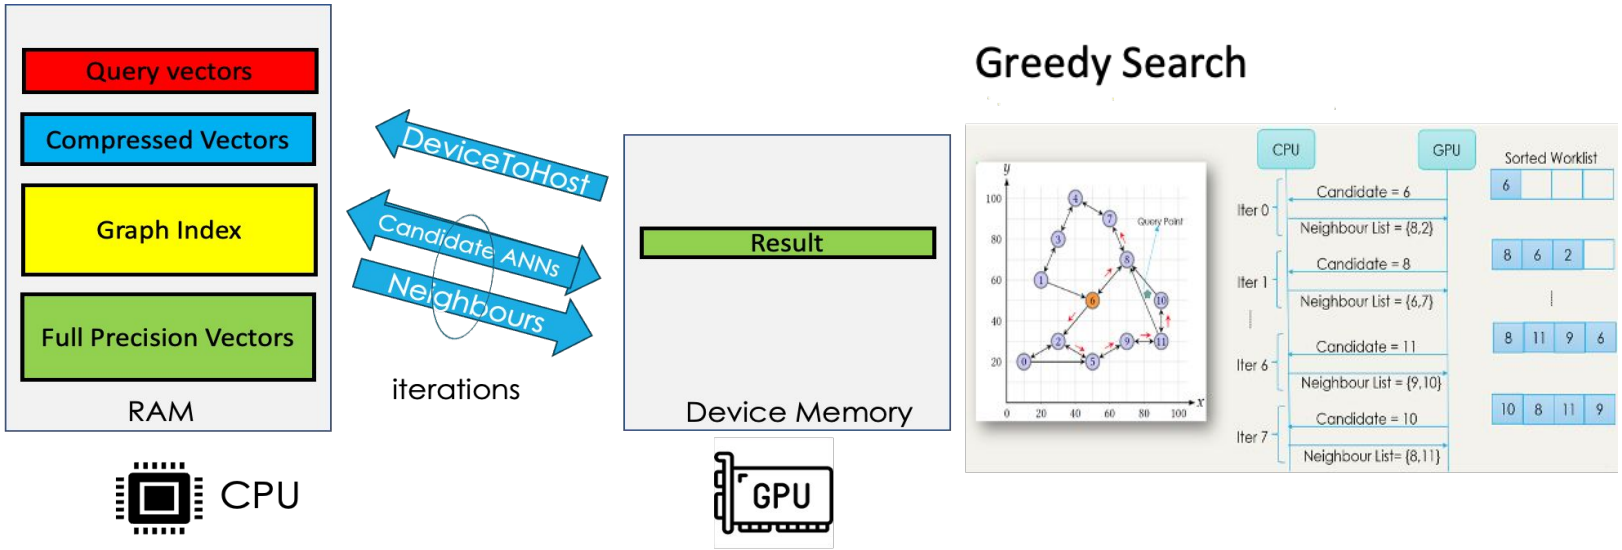

Program Analysis using LLMs

- Does CodeT5, PLBART, CodeBERT, UniXcoder, etc., perform semantic tasks well (e.g., bug prediction, bug fixing, etc.)?
- Cell2Doc: ML Pipeline for Generating Documentation in Computational Notebooks [**ASE'23**]

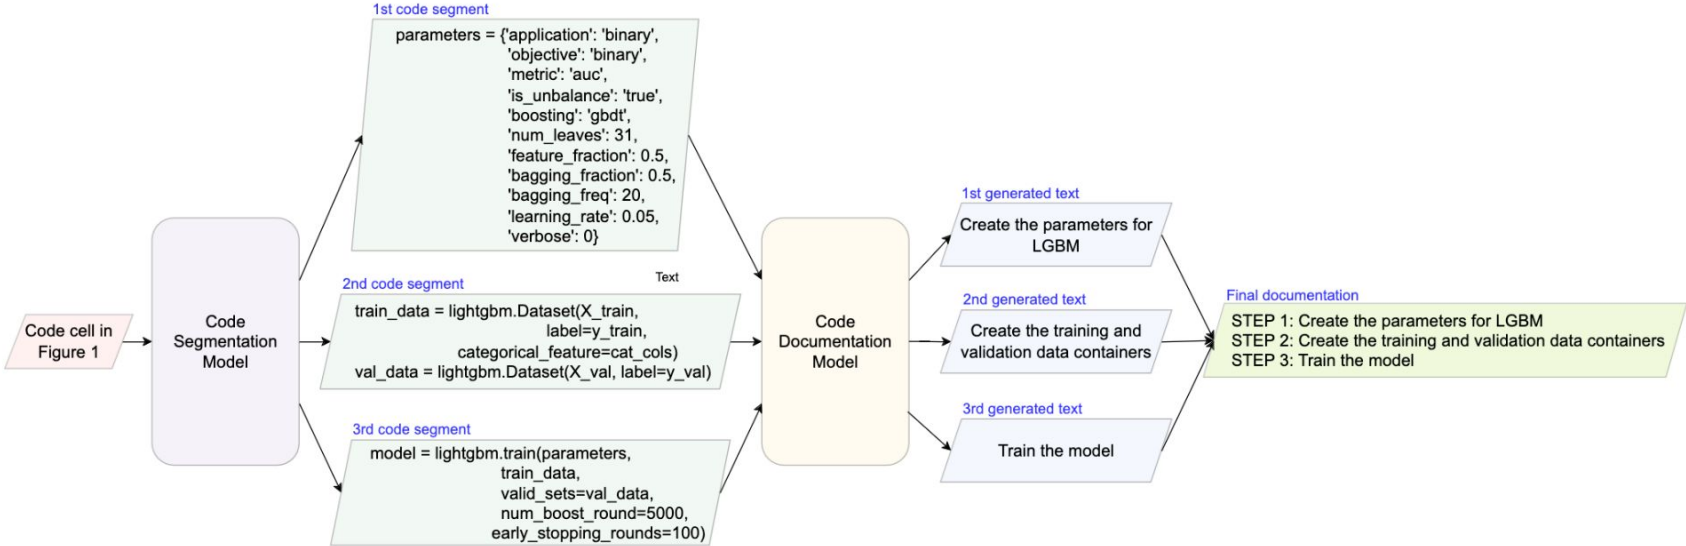

Programming Languages & Synthesis: Broad Classification of our Research Efforts :

Analysis and Refactoring

Automated Verification, Rich Type Systems

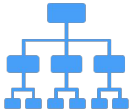

Symbolic Program Synthesis

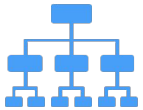

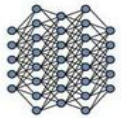

Neurosymbolic Synthesis

Analysis

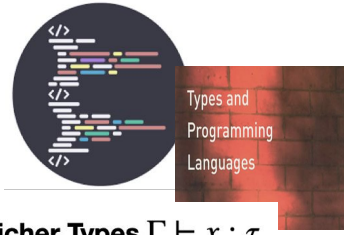

Richer Types $\Gamma \vdash x : \tau$

Spec:

$$\Psi = x : \tau \rightarrow \{P\}v : t\{Q\}$$

FM4Synthesis

Spec: i/o examples, A NL Query

$$\Psi = x : \tau \rightarrow \{P\}v : t\{Q\}$$

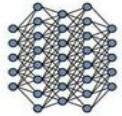

FM4Repair

[ESME 22, ICSE 23]

ECOOP' 23, PLDI 23

OOPSLA 22

Ongoing Efforts

Artificial Intelligence/Machine Learning Research @IITH

Deep Learning Architecture and Training

- Explore
  - New architectures and models
  - New training methods and loss functions
  - Newer inputs

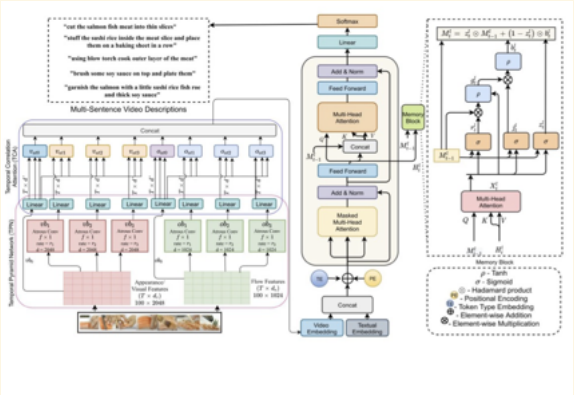

Generative AI

- AI to create a wide variety of data, such as images, videos, audio, text and 3D models
- GenAI learns patterns from existing data and uses that knowledge to generate new and unique data.
  - GenAI applications: ChatGPT, DeepBrain, Synthesia,...

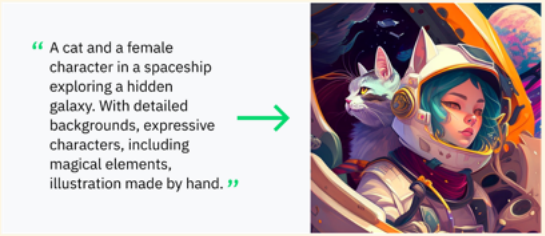

Image Generation

Image Courtesy: generativeai.net

Bayesian Learning

- Use Bayesian Learning for building Safe AI applications
- High risk real-world applications, such as autonomous vehicles and healthcare
- Bayesian Deep Learning provides better decision making by handling uncertainty, robustness and considering domain knowledge

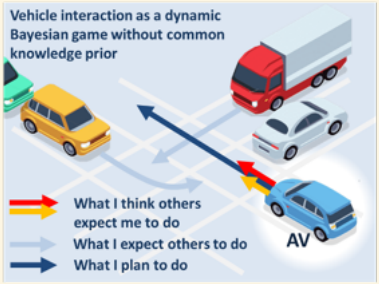

Algorithms for Massive Datasets

- Developing algorithms for handling large dimensionality and large volume of datasets
  - High-dimensional: text vocabulary, pixels in image
  - Large volume: millions of documents and images
- Develop efficient distributed algorithms – Hadoop/MapReduce
- Use Sketching/Sampling to turn “Big Data into tiny data”

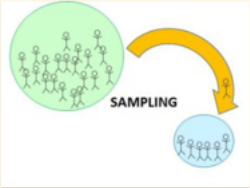

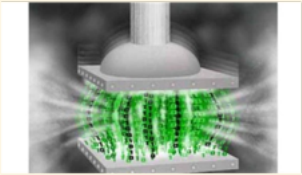

Computer Vision

- Action recognition, emotion recognition and video analytics
- Autonomous vehicle technology
- Aerial imagery analysis and image captioning
- Medical imaging

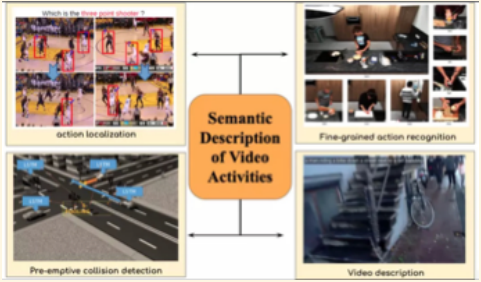

Vision for Drones

- Challenge: Handling objects of different sizes
- Applications: Surveillance, Search and Rescue, Infrastructure Inspection, Crop Health Monitoring, Land Cover Mapping, Traffic Management
- Detecting drones from drones

NLP: Dialog Systems

- Computer system intended to converse with a human.
- Uses one or more of text, speech, graphics, gestures, etc., to communicate between human and system
- Build scalable and explainable dialog systems

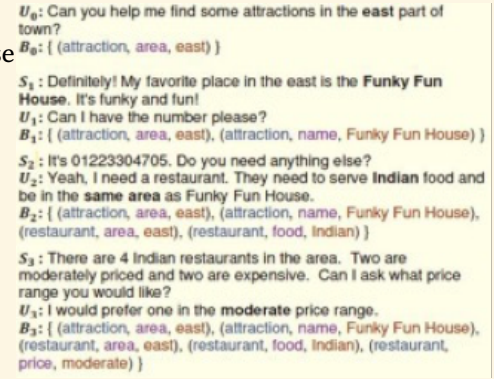

NLP: Personalized Autosuggest

- Personalized query autocompletion for short and unseen prefixes.

Let previous ***n*** queries (earliest to latest order) in the current **session s** be $\{q_1, q_2, \ldots, q_n\}$. Current **query** is **q**, and **p** is the **query prefix** typed so far.

Generate top-**N** query completions conditioned on current **query prefix p**, additional trie context **e**, and session information **s** i.e., $P_\theta(q \mid p; e; s)$

Social Media Analytics

- Information diffusion in social network
- Summarize social media content
- Categorize content
  - Spam vs non-spam
  - Quality of posts and replies
- Content routing

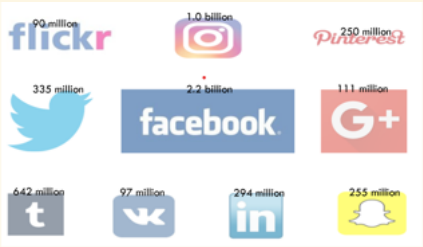

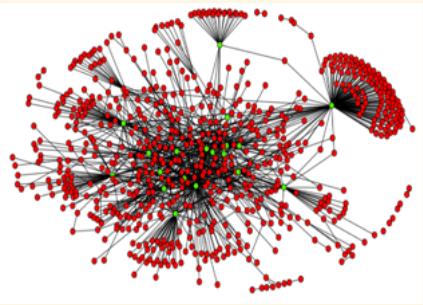

Applied Learning Theory

- Kernel methods
- Statistical learning theory
- Optimization
- Generative AI

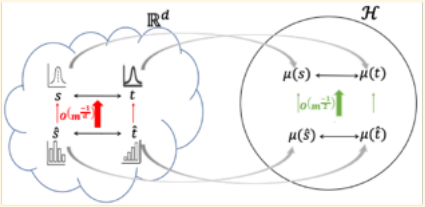

Machine Learning Faculty

C Krishna Mohan

Video Content Analysis, Machine Learning, Sparsity Based Methods, Deep Learning

Manish Singh

Databases, Data Mining, Information Retrieval

Sobhan Babu

Big Data Analytics, Graph Theory and Applied Algorithms

Srijith P.K.

Bayesian Data Analysis, Probabilistic Machine Learning, Survival Analysis and Text Analytics

Saketha Nath Jagarlapudi

Machine Learning

Maunendra Desarkar

Recommender Systems, Information Retrieval

Vineeth N. Balasubramanian

Machine Learning, Computer Vision

CSE PhD Alumni

Collaborations

Publication Venues

Infrastructure

- Large number of servers (with CPU and GPU) available through SLURM and MAAS infrastructure
- High Performance Computing (HPC) cluster available under the National Supercomputing Mission
- OpenStack-based private cloud for Virtual Machines (VMs)
- A large range of IoT and embedded processors and FPGA boards
- State-of-the-art DGX Servers
- High-end network switches such as 3.2 Tbps Intel Tofino Programmable Switch
- Labs with dedicated workstation for every Ph.D. scholar, with 24x7 access

Fellowships

- Ministry of Education (MoE) fellowships
- Sponsored research project fellowships
- Joint Ph.D. fellowships with IDBRT Hyderabad, Swinburne University Australia, and Deakin University Australia
- Industry fellowships such as Google, TCS, Intel fellowships
- PM Research Fellowship (PMRF)
- Visvesvaraya Fellowship
- Financial assistance for Ph.D. students to present their research papers in international and national venues

Life @CSE, IITH

- PhD seminar talks:
  CSE Ph.D Seminar Talks IIT-Hyderabad
- Several international and national computer science conferences held @IITH, e.g. ACML 2022, CALDAM 2020, etc.
- Research Scholars Day