

Usage of SERC SahasraT Supercomputer for Scientific Applications in IISc in 2019



Supercomputer Education and Research Centre

Indian Institute of Science

Bangalore – 560012

May 2020

Table of Contents

1. SERC Cray SahasraT System and NVIDIA DGX-1 System.....	3
2. Usage in 2019	3
3. IISc's Scientific Applications on SahasraT and DGX-1	7
3.1 Department of Aerospace Engineering (AE)	8
3.2 Centre for Brain Research (CBR)	13
3.3 Department of Computational and Data Sciences (CDS)	14
3.4 Department of Chemical Engineering (CE).....	19
3.5 Centre for Earth Sciences (CEaS)	21
3.6 Department of Electronic Systems Engineering (DESE).....	22
3.7 Inorganic and Physical Chemistry (IPC).....	23
3.8 Materials Research Centre (MRC).....	25
3.9 Department of Physics (PHY).....	29
4. Access to External Users	30
5. Publications in 2019 using SahasraT.....	30
6. Projects Based on SahasraT.....	33

1. SERC Cray SahasraT System and NVIDIA DGX-1 System

SERC's current supercomputing resources include a 33000+ CPU-core 1.1 PetaFlop Cray XC40 supercomputer for traditional HPC simulations and a NVIDIA DGX-1 for machine learning and AI.

SERC is home to one of the fastest supercomputers in India, the Cray XC40, known as **SahasraT** in IISc. It was commissioned in 2015. SahasraT is also among the fastest academic supercomputers in India. It is a system that combines the capabilities of Intel's latest Xeon Haswell processors for the CPU cluster and Nvidia's K40 series of GPU cards and Intel's energy efficient Xeon Phi Knights Landing MIC processors as a cluster connected using Cray's own Aries high-speed interconnect on a dragonfly topology with DDN's high performance storage units.

It is a 1.4 PFlop system with 33000 Intel Haswell cores with total memory of 176 TB RAM, and 2 PB of storage. It also has 44 accelerator nodes each with one NVIDIA K40 GPU card, and 24 nodes each with one Intel Xeon KNL processors.

System also hosts architecture specific parallel libraries like OpenMP, MPI, CUDA and Intel Cluster software. Extensive range of parallel Scientific and Mathematical libraries like BLAS, LAPACK, Scalapack, fftw, hdf5, netcdf, PETSc, Trilinos etc. are also available on the system.

The CRAY XC40 facility uses PBS (Portable Batch System) to schedule jobs. The scheduler for batch jobs in SahasraT uses the FCFS (First-Come, First-Served), based on queue priority and backfilling techniques to achieve the highest system utilization possible with a reasonable turn-around time.

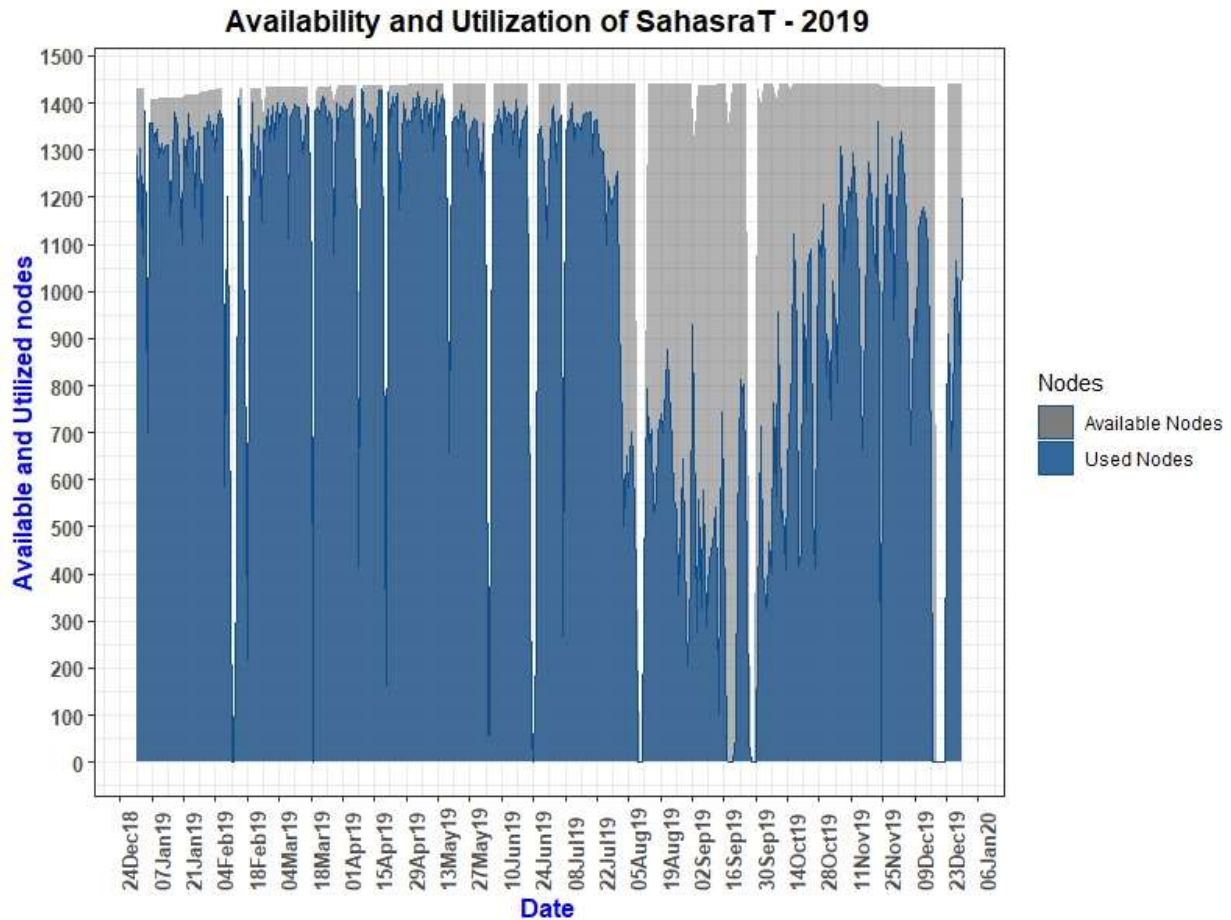
The **NVIDIA DGX-1** is a deep learning system, architected for high throughput and high interconnect bandwidth to maximize neural network training performance. The core of the system is a complex of **Eight Tesla V100 GPUs** connected in the hybrid cube-mesh NVLink network topology. In addition to the eight GPUs, DGX-1 includes two CPUs for boot, storage management, and deep learning framework coordination. DGX-1 is built into a three-rack-unit (3U) enclosure that provides power, cooling, network, multi-system interconnect, and SSD file system cache, balanced to optimize throughput and deep learning training time.

NVLink is an energy-efficient, high-bandwidth interconnect that enables NVIDIA GPUs to connect to peer GPUs or other devices within a node at an aggregate **bi-directional bandwidth of up to 300 GB/s per GPU**: over nine times that of current PCIe Gen3 x16 interconnections. The NVLink interconnect and the DGX-1 architecture's hybrid cube-mesh GPU network topology enables the highest achievable data-exchange bandwidth between a group of eight Tesla V100 GPUs.

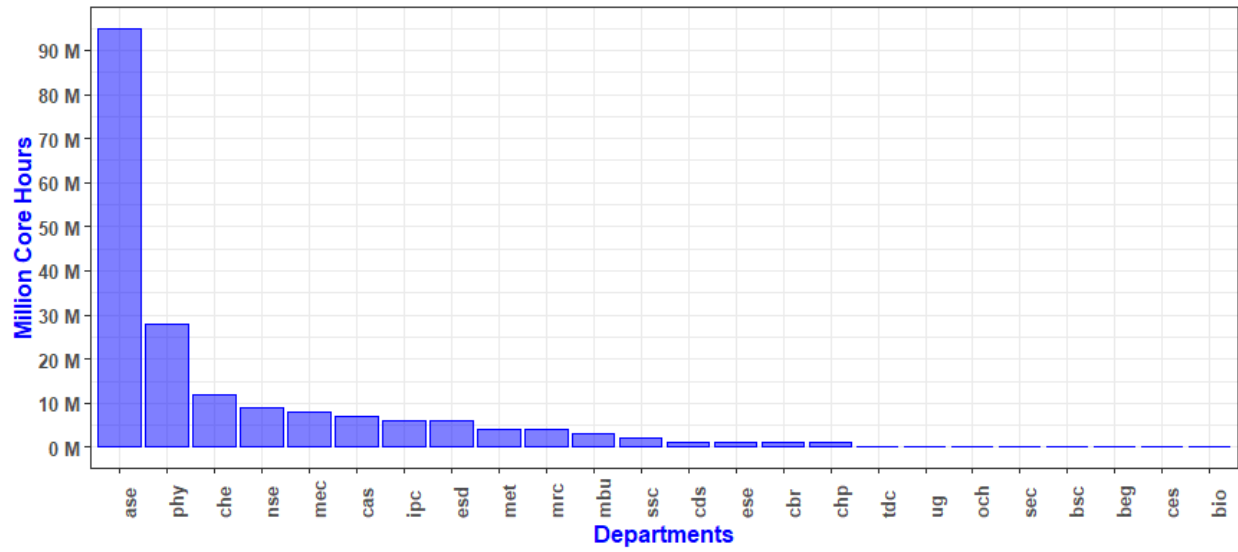
2. Usage in 2019

SahasraT serves 44 Departments, 134 Research Labs and over 450 users. Following are some salient statistics of the usage of SahasraT in 2019.

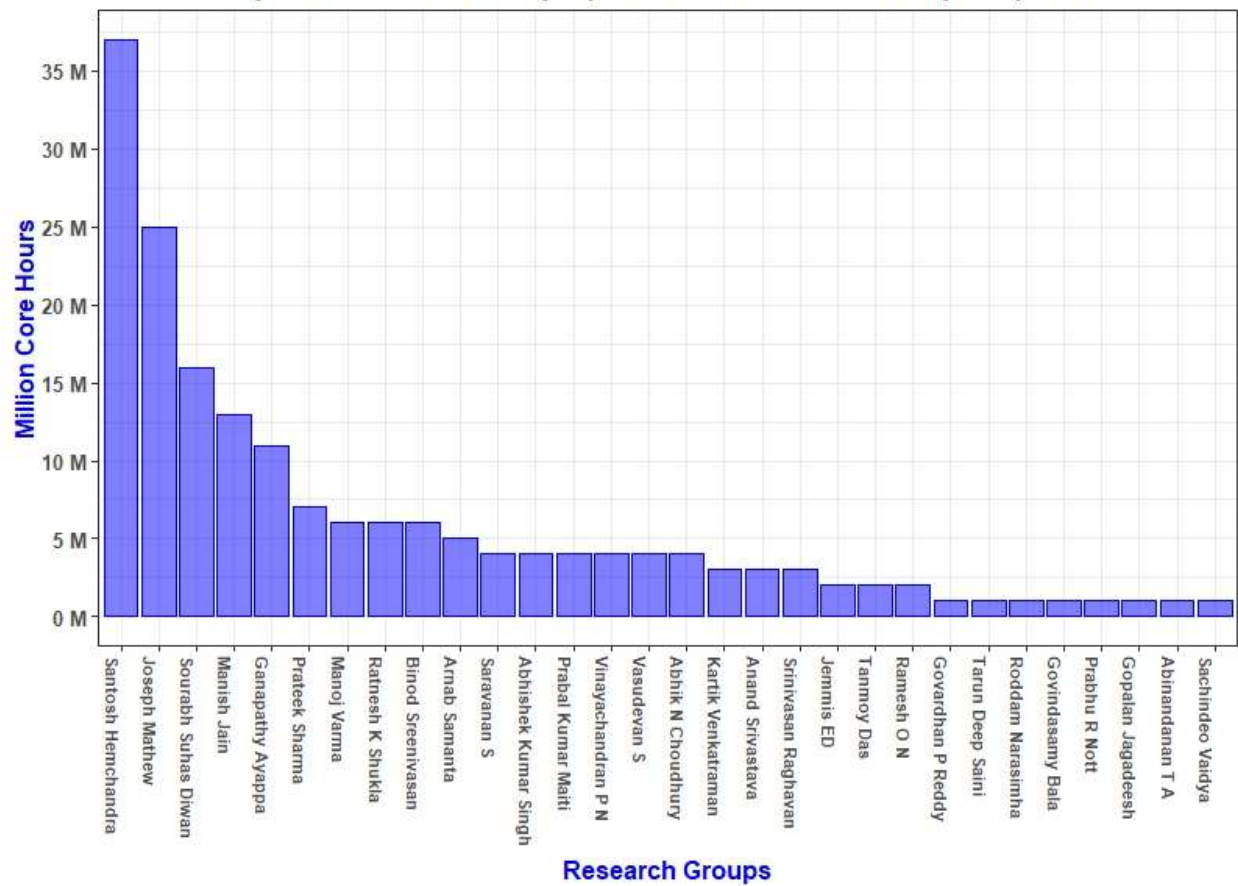
- A total of 126400 jobs were processed in 2019, at an average of 10533 jobs per month. 74.5% of the jobs were CPU jobs and 23.5% of the jobs were GPU jobs, the remaining were jobs in in special queues like temporary queue or advance reservation queues.
- The total system utilization was about 176 Million core hours i.e., an average system utilization of about 14.7 million core hours per month.
- The overall availability and utilization for 2019 and usage by the different departments in 2019 is given in the following charts.




Core Hour Utilization by Department - 2019



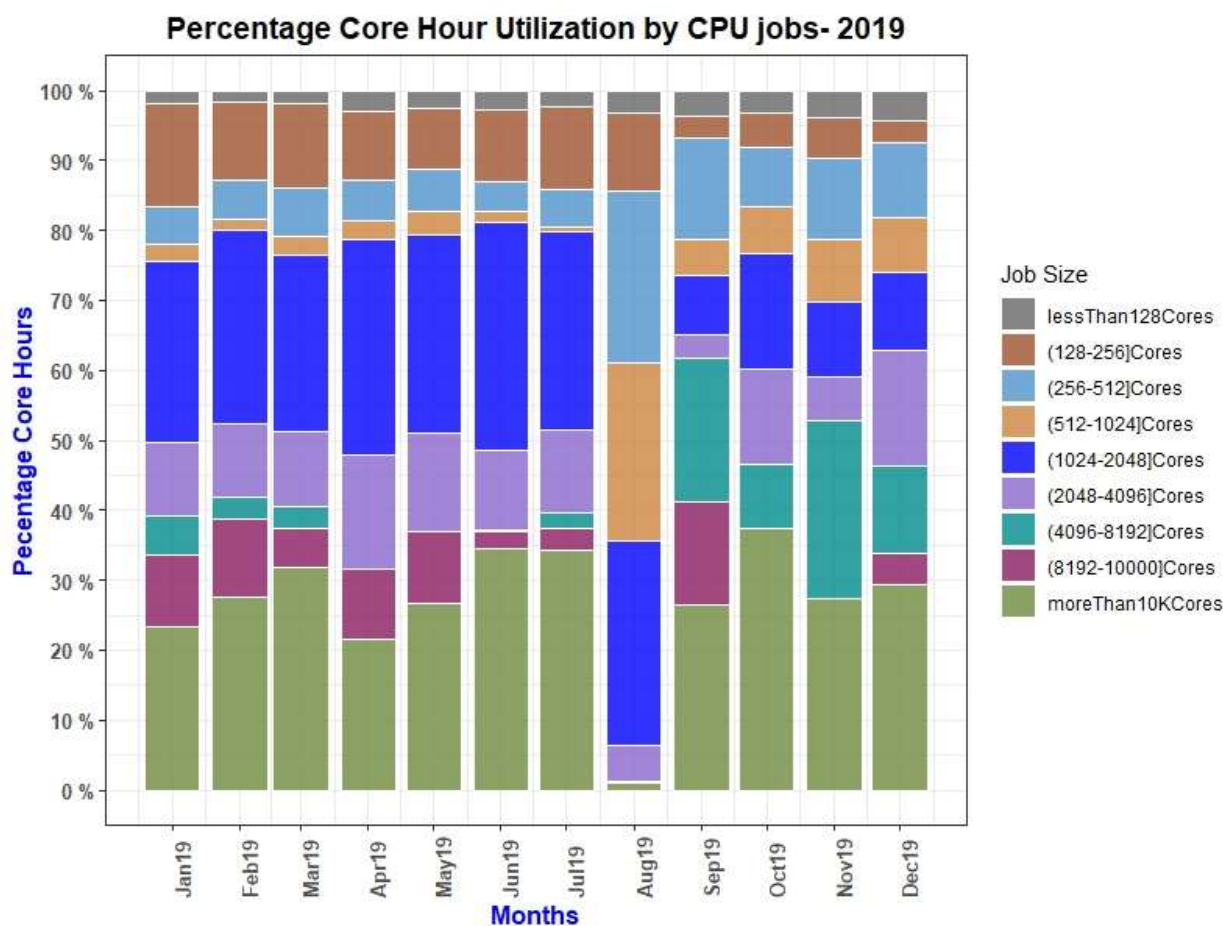
Top 30 Research Groups (CPU Core Hour Consumption) - 2019



The following chart provides resource utilization in terms of core hours per regular CPU Queue.



561 large jobs were processed in 2019 as compared to 667 in 2018 and only 26 large jobs in 2017.



26 percent of the CPU hours were consumed by jobs using 1024 to 2048 cores, 27 percent of CPU hours were consumed by capability jobs using more than 10000 Cores and 7 percent of CPU hours were consumed by capability jobs using more than 8192 Cores.

3. IISc’s Scientific Applications on SahasraT and DGX-1

In 2019, a large number of scientific applications were executed by the Institute community on SahasraT. These belonged to different departments including Aerospace, Physics, Centre for Atmospheric and Oceanic Sciences (CAOS), Centre for Brain Research (CBR), Department of Computational and Data Sciences (CDS), Chemical Engineering (CE), Centre for Earth Sciences (CEaS), Department of Computer Science and Automation (CSA), Department of Electronics Structure and Engineering (DESE), Inorganic and Physical Chemistry (IPC), Materials Engineering, Materials Research Centre (MRC), Molecular Biophysics Unit (MBU) and Mechanical Engineering. The applications spanned different domains including

- aerodynamics, turbulence computations,
- climate modelling,
- whole genome sequencing,
- study of proteins in breast cancer,

- large-scale finite element methods,
- large-scale graph problems,
- study of granular materials,
- Bayesian optimization methods,
- 2-D materials for lithium ion battery,
- interactions in halogen bond, carbon nanotubes,
- microstructure,
- study of biological membranes, analysis of conformational space of proteins,
- computational modelling of materials using machine learning,
- thermal transport in glass-forming liquids,
- study of band insulator,
- study of bilayer transition metal dichalcogenides (TMDs),
- study of biomolecules, glycoproteins, gene silencing mechanisms, dendrimer, DNA nanotubes,
- study of superfluid turbulence,
- etc.

The following sections give details of the usage of SaharsaT for various computational science problems.

3.1 Department of Aerospace Engineering (AE)

3.1.1 Prof. Santosh Hemchandra's Lab

3.1.1.1 *Research*

Unsteady flow dynamics in a gas turbine swirl nozzle

This study aims to clarify the impact of nozzle centerbodies on the unsteady flow dynamics of swirl nozzles. These flow dynamics can impact pollutant formation levels and the thermoacoustic characteristics of the combustor. Therefore, controlling these flow instabilities by passive, nozzle design modifications is desirable. We focus our efforts on a swirl nozzle combustor whose flow field has been characterized experimentally by the Lieuwen group at Georgia Tech. Our study shows using a combination of Large eddy simulations and linear stability analysis methods, that the unsteady helical breakdown bubble precession mode is activated when the vortex breakdown bubble generated by swirl and the wake generated by the centerbody are spatially separated.

3.1.1.2 *Parallelization*

Our solver uses a domain decomposition strategy to solve the Navier-Stokes equations using the finite difference method.

3.1.1.3 *SERC Resources and Experiments*

Computations have been performed using 12,000 – 15,000 cores per run on the SaharsaT Cray XC40 CPU system. The mesh size is between 80-90 million points.

3.1.1.4 *Performance*

The cases run for this project cannot be run on our lab cluster. However, on cases that can be run on both CRAY and lab cluster, we see a 1.5x speedup on the cray relative to the lab cluster. The latter has the same compute cores as CRAY (Intel Haswell) but uses commodity FDR Infiniband for the backplane interconnect.

3.1.1.5 Publications

Mukherjee, A., Muthichur, N., More, C., Gupta, S. and Hemchandra, S., “The role of the centerbody wake on the precessing vortex core dynamics of a swirl nozzle”, Journal of engineering in gas turbines and power, *to appear*.

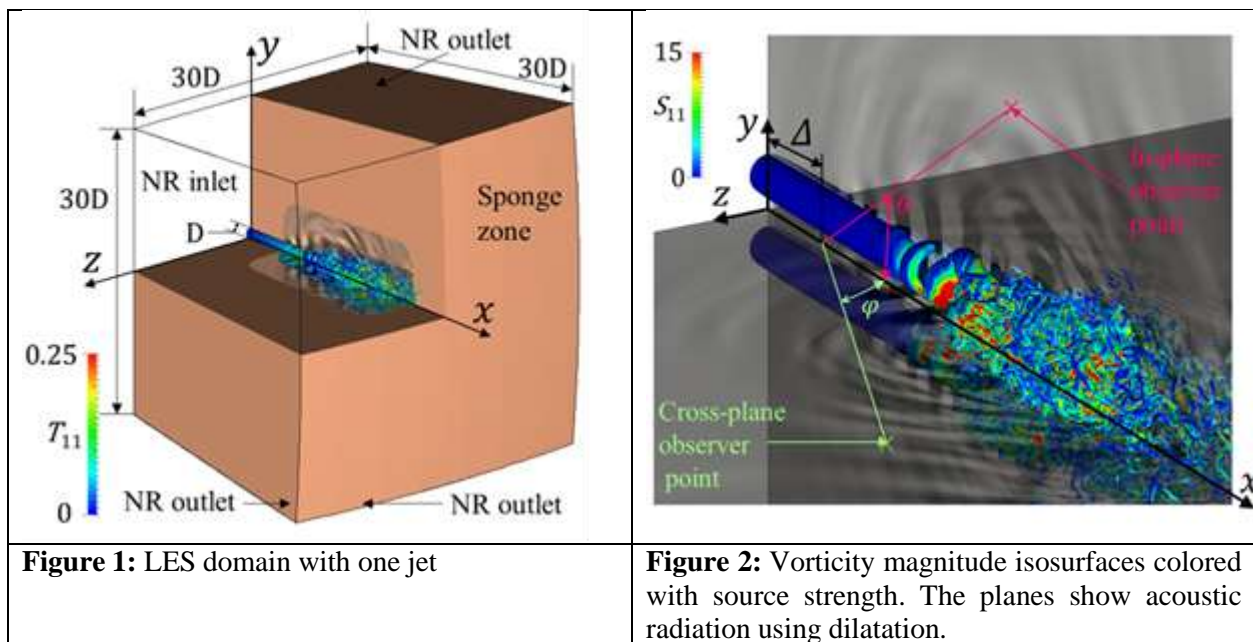
3.1.1.6 Impact

This study provides a preliminary design guideline for combustor nozzle designers on how helical flow oscillations that help improve mixing and thermoacoustic characteristics in swirl nozzles may be induced. This work has generated interest in collaborative projects from several OEMs.

3.1.1.7 Research

Acoustic radiation from twin M=0.9 jets

Most aerospace vehicles such as rockets and aircraft feature high momentum and temperature jets exiting the vehicle. These jets produce high levels of acoustic radiation over significantly a wide frequency range that can adversely impact the structural integrity of structures around launchpads and runways, as well as, detrimentally impact the health of personnel working routinely around these vehicles. The present study aims to understand the spectral characteristics of acoustic radiation in the far-field of twin round jets at an exit Mach number of 0.9 by varying the inter-jet spacing.



3.1.1.8 SERC Resources and Experiments

Computations have been performed using 12,000 cores per run on the SaharaT Cray XC40 CPU system. The mesh size is about 96 million points.

3.1.1.9 Parallelization

Our solver uses a domain decomposition strategy to solve the Navier-Stokes equations using the finite difference method.

3.1.1.10 Performance

The cases run for this project cannot be run on our lab cluster. However, on cases that can be run on both CRAY and lab cluster, we see a 1.5x speedup on the cray relative to the lab cluster. The latter has the same compute cores as CRAY (Intel Haswell) but uses commodity FDR Infiniband for the backplane interconnect.

3.1.1.11 Publications

Muthichur, N., Harikrishna, T., Hemchandra, S. and Samanta, A., “Sources of Sound and its Radiation from Twin Turbulent Jets”, AIAA aerospace science and technology forum, 2020, AIAA paper# 2020-1245.

3.1.1.12 Research

Detailed quantitatively and qualitatively accurate modelling of reacting turbulent flows in aerospace systems remains a significant challenge due to the fact that the interaction between the flame and turbulent flow scales is a significant process that must be correctly captured by the simulation. We explore a new approach to modelling reacting flows using the explicit filtering large eddy simulation (EFLES) method. This work compares results from DNS with corresponding LES simulations of a turbulent premixed methane-air jet flame in the thin reaction zones regime.

3.1.1.13 SERC Resources and Experiments

Computations have been performed using 12,000 cores per run on the SaharaT Cray XC40 CPU system. The mesh size is about 96 million points.

3.1.1.14 Parallelization

Our solver uses a domain decomposition strategy to solve the Navier-Stokes equations using the finite difference method.

3.1.1.15 Performance

The cases run for this project cannot be run on our lab cluster. However, on cases that can be run on both CRAY and lab cluster, we see a 1.5x speedup on the cray relative to the lab cluster. The latter has the same compute cores as CRAY (Intel Haswell) but uses commodity FDR Infiniband for the backplane interconnect.

3.1.1.16 Publications

Datta, A., Mathew, J., and Hemchandra, S., “The explicit filtering method for large eddy simulations of a turbulent premixed flame”, submitted to the International symposium on combustion, 2020 (*under review*)

3.1.1.17 Impact

The EFLES technique has been shown to be effective in capturing the qualitative and quantitative features of complex non-reacting turbulent flows in several studies. The method models the impact of turbulence processes at the filter cut-off scale and shows very good quantitative accuracy characteristics for a given mesh size. The hope is that extending this modelling approach to reacting flow, will significantly impact the way LES is used by OEMS to design low pollutant emissions combustion systems.

3.1.2 Prof. Sourabh Suhas Diwan's Lab

Team: Transition and Turbulence Research Group

3.1.2.1 Research

1. Turbulent Entrainment in steady jets and plumes with and without heat addition.
2. Transition and Intermittency in a Laminar Separation Bubble

1. A study seeking entrainment behavior of free shear flows such as a jet or plume and the effect of heat addition has been conducted (figure 1a). Jet and plume flows have been established in lab with validation of their first and second order statistics (figure 1b). Similarity in vorticity fluxes of different free shear flows (jet & plume) has been found and this has been reported at the recently held Asian Congress of Fluid Mechanics in December 2019 in Bangalore. Attempts are being made to understand the behavior of vorticity fluxes in relation to the coherent structures in the flow (figure 1c).

Figure 1 (a) Passive scalar contour at value of 0.01 of a turbulent jet at Reynolds number 3200, with a sub-image of plume of Lockheed rocket test fired in Los Angeles hills ($Re \approx 108$), indicating similarity of form.

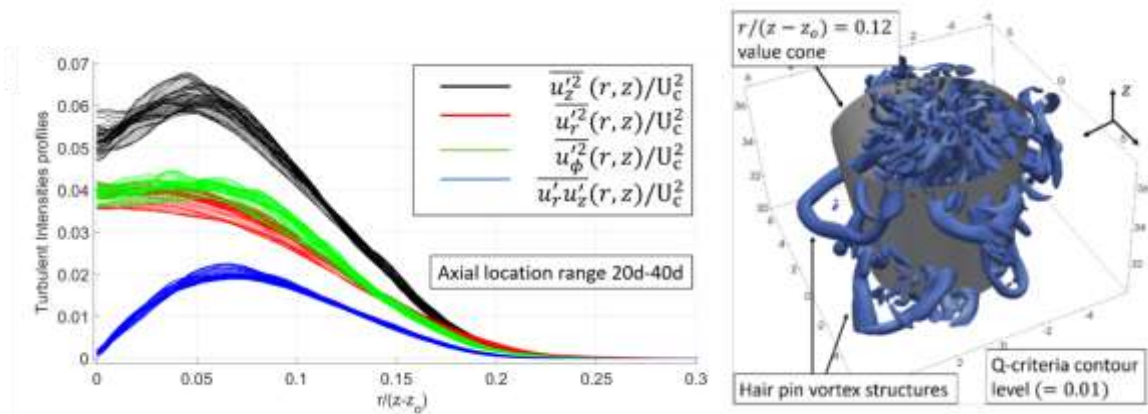
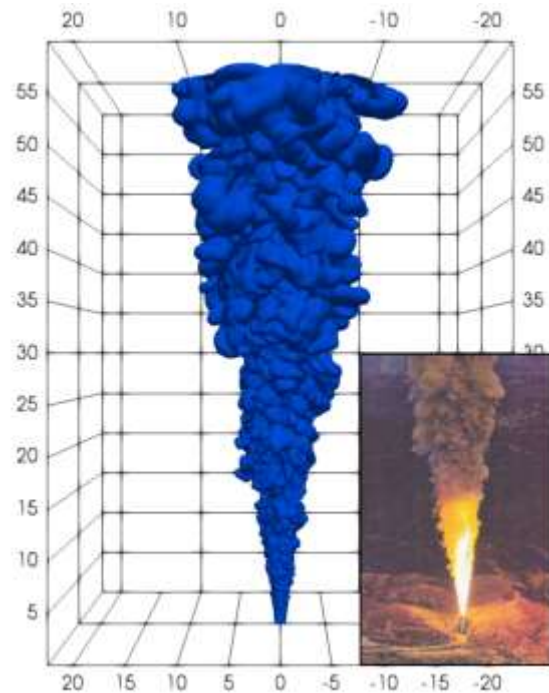


Figure 1 (b) Radial profiles of the second-order turbulent statistics of a jet. (c) Iso-surface of the Q-criterion for plume representing hair-pin structure in outer region.

2. A study investigating the change in transition mechanisms in flows under the influence of adverse pressure gradient has been carried out. Study of intermittency & turbulent spots in separated shear layer and their comparison with those in an attached boundary layer is being pursued. Figure 2 shows the way spanwise vortical rollers are distorted and broken up due to their interaction with streaks generated upstream and leading to the generation of significant 3D motions.

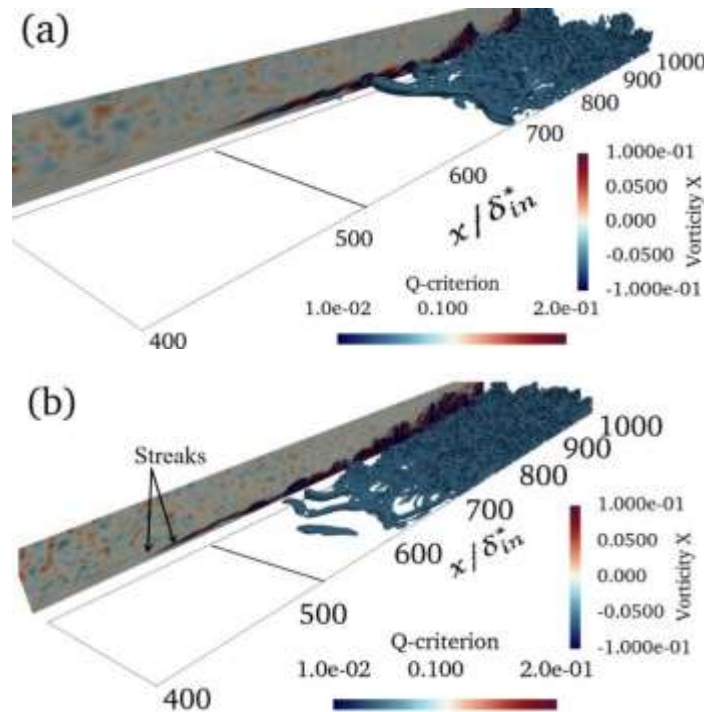


Figure 2: Interactions of streaks with separated boundary layer resulting in change in transition mechanism at two Reynolds numbers: (a) $Re_{\square^* in} = 79.2$ and (b) $Re_{\square^* in} = 105.8$.

3.1.2.2 SERC Resources and Experiments

13.835 million hours of SahasraT were used.

3.1.2.3 Parallelization

An Incompressible code parallelized using MPI. Pencil fast Fourier transform with “fftw” in background has been incorporated to solve Poisson equation.

3.1.2.4 Performance

Turbulent jet flows were simulated using 2916 cores (122 nodes) of Cray XC40. Other smaller simulation runs also performed. No machine available in lab for comparison.

3.1.2.5 Publications

Rohit Singhal, S. Ravichandran, Sourabh S. Diwan and Garry L. Brown, *Reynolds stress gradient and vorticity fluxes in axisymmetric turbulent jet and plume*, The 16th Asian Congress of Fluid Mechanics, 13-17 December 2019, JNCASR Jakkur Bengaluru, India.

3.1.2.6 Other Outcomes

MTech (course) Student Thesis: Kartik Naicker, 2019 (Aerospace Engg. Dept): *Direct numerical simulation of transition in a laminar separation bubble.*

3.2 Centre for Brain Research (CBR)


3.2.1 Prof. Bratati Kahali's Lab

3.2.1.1 Research

Population Level Genomic Variant Calling

Research Statement: To identify large population level genomic variations with at least thousands of human whole genome sequencing data. This has been achieved with the help of datastores that are efficient in processing the variants in a single process multiple data (SPMD) fashion which leads to increased performance in fetching the genomic data for tens of thousands of human samples.

Experimental Setup: 600 cores of Sahasrat Cray XC40 were used for this experiment. The jobs were submitted in small queue upon 25 nodes. Total of 53 samples g.vcf files were consolidated into chromosome wise datastores using special genomics datastore utilities of GATK in 25 individual nodes of Cray. Upon the generated datastore using genotypeGVCF of GATK joint call is performed and final chromosome level VCF files are obtained. MPI based parallel codes in python were written which helped in processing 25 chromosomes [chromosome-1 to 22; X, Y, M] in parallel upon 25 compute nodes (each chromosome mapped to each node).



3.2.1.2 SERC Resources and Experiments

SahasraT Usage: 31,800 core hours.

3.2.1.3 Performance

The total size of all the datastores generated is 750GB. The total size of project level vcf consisting up of 53 samples is around 25 GB. Utilizing 25 nodes in parallel in Sahasrat small72 queue helped in completing the joint variant calls from 53 human samples in 52.74 hours. The maximum wall-time was 53 hours for

each jobs associated per chromosome of pipeline. Each chromosome of the joint calling pipeline has different completion time which is depicted below. The longest time is taken by chromosome 2 which is around 52.74 hours and the smallest time taken is by chromosome M which took 10.8 seconds.

3.2.1.4 Research

Genome Wide Genetic Interaction Analysis.

Research Statement: To identify interacting loci using genotype matrices and corresponding population demographics data

Experimental Setup: 960 cores of Sahasrat Cray XC40 were used for this experiment. The jobs were submitted in small queue upon 40 nodes; each node utilizing 24 cores. MPI based parallel codes in python were written to create randomized iterations which helped in processing 40 instances simultaneously upon 40 compute nodes (1 randomised instance mapped to 1 node). MDR is the method used to generate interacting loci models. The preprocessing stages involved randomization and reduction of loci distribution. The mentioned software was parallelized in a multi-core fashion within the nodes using the available internal thread options.

3.2.1.5 Performance

The input size of the genotype matrix comprised of 596 subjects marked 5279 genomic loci. Utilizing 40 nodes in parallel in Sahasrat small queue helped in identifying the interactions in the current setting within 10 hours of walltime.

3.3 Department of Computational and Data Sciences (CDS)

3.3.1 Prof. Partha Talukdar's Lab

3.3.1.1 Research

Extreme classification on Graphs

The lab worked on a specific case of node classification on graph data, where the number of labels is very large (~500K-600K). Existing methods (e.g., GCNs) may not scale very well due to large number of parameters at the final output (softmax) layer. The lab explored and experimented with various strategies to address this issue. *Application:* This problem appears in web applications, like product recommendation, wikipedia page category/tag predictions, etc.

.

3.3.1.2 SERC Resources and Experiments

NVIDIA-DGX. Two GPUs were used.

3.3.1.3 Performance

Due to the large number of labels and nodes in the graph dataset, the model size exceeded the lab GPU RAM capacity (~12GB). In NVIDIA-DGX, the model took ~8-9 hours to run.

3.3.2 Prof. Sathish Vadhiyar's Lab

3.3.2.1 Research

Multi-node Multi-device Louvain Algorithm for Community Detection

Anwasha Bhowmick, Sathish Vadhiyar

This project proposed a novel hybrid CPU-GPU algorithm for large-scale community detection in graphs. A divide-and-conquer strategy was followed where communities are formed on different parts of the graph simultaneously on CPU and GPU. Subsequently, the doubtful vertices are resolved during the merging step. The work was extended to multiple nodes of CPU and GPU, where a novel merging step is employed.

3.3.2.2 Parallelization

The work employed a novel divide-and-conquer paradigm for simultaneous formation of communities on the different parts of the graphs on the CPU and GPU cores.

3.3.2.3 SERC Resources and Experiments

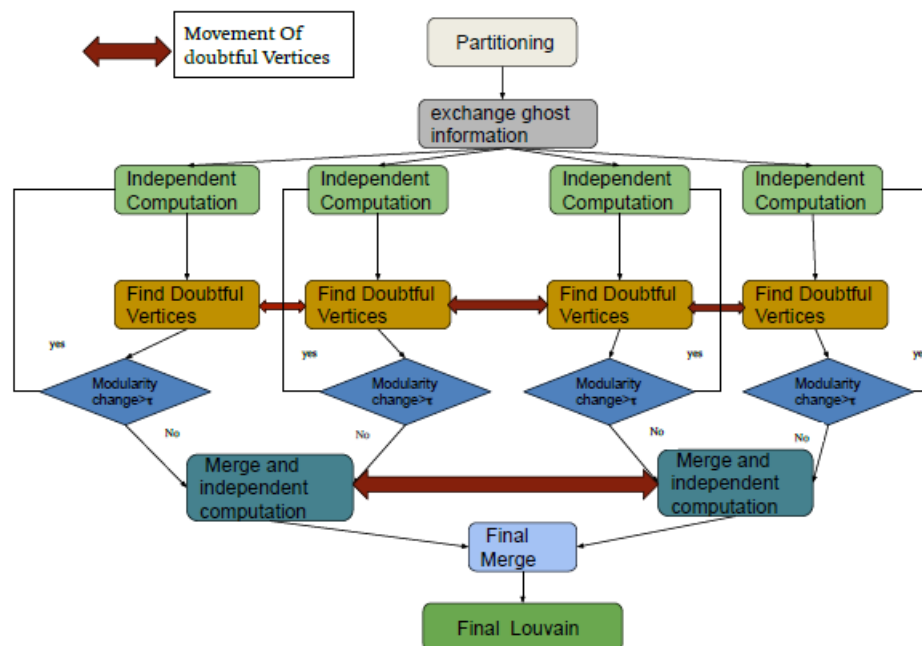
SahasraT GPU nodes were used for this purpose. Experiments were performed up to 8 and 16 nodes of GPUs, where all the CPU and GPU cores were utilized.

3.3.2.4 Performance and Scale

SahasraT provides a multi-GPU node environment that helps explore billion-vertex problems that can otherwise not be explored on lab machines. Our experiments with multiple nodes show good scalability of the multi-node multi-GPU algorithm with increasing number of nodes.

3.3.2.5 Publications

Anwasha Bhowmik, Sathish Vadhiyar: HyDetect: A Hybrid CPU-GPU Algorithm for Community Detection. HiPC 2019: 2-11.



3.3.2.6 Research

Fast and Accurate Learning of Knowledge Graph Embeddings at Scale

Udit Gupta, Sathish Vadhiyar

Knowledge Graph Embedding (KGE) is used to represent the entities and relations of a KG in a low dimensional vector space. KGE can then be used in a downstream task such as entity classification, link prediction and knowledge base completion. This work proposed three strategies which lead to faster training in a distributed setting. Combining the three strategies results in reduction of training time for the FB250K dataset from twenty-seven hours on one processing node to under one hour on thirty-two nodes with each node consisting of twenty-four cores.

3.3.2.7 Parallelization

Data parallelism was employed across the distributed nodes. Novel algorithms were proposed to minimize the size of the allgathers and for distributed memory Adam optimization approach.

3.3.2.8 SERC Resources and Experiments

SahasraT CPU nodes were used for this research. Up to 32 nodes (1536 cores) were used for this purpose.

3.3.2.9 Performance and Scale

The large number of nodes helped us to explore large-scale graph embedding problems which was not possible in our lab cluster.

3.3.2.10 Publications

Udit Gupta, Sathish Vadhiyar: Fast and Accurate Learning of Knowledge Graph Embeddings at Scale. HiPC 2019: 173-182.

3.3.2.11 Research

Scalability Bottleneck Analysis of High Performance Applications

Soham Halder, Sathish Vadhiyar

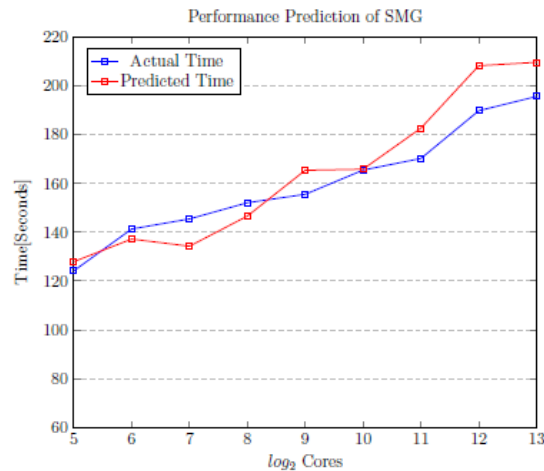
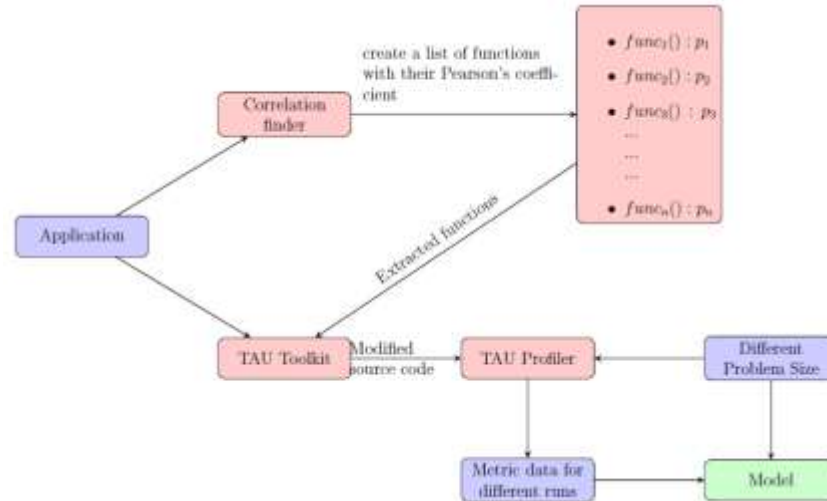
Prediction of scalability trends of high performance applications for large scale executions using small number of cores is important and has practical uses including helping the developers tune their codes for future large-scale systems, helping supercomputing centers to plan future large-scale deployments etc. In this work, we develop strategies for both performance predictions and scalability projections including for applications whose scalability trends can change at scale due to the presence of scalability bottlenecks. The work uses strategies including small-scale benchmarks for prediction of behaviour of large-scale applications and using number of function calls for time consuming functions to predict scalability trends. The techniques were demonstrated for scientific applications including GTC, MILC, SMG, LAMPPS and AMG.

3.3.2.12 SERC Resources and Experiments

SahasraT CPU resources were used for this research. Up to 16000 core executions were made.

3.3.2.13 Performance and Scale

Such large-scale executions are not possible to be carried out in the department and lab clusters.



3.3.2.14 Research

Combining Approaches of Compressing Neural Networks

Mayank Gupta, Sathish Vadhiyar

Neural networks are the model of choice for many problems in computer vision, Natural Language Processing, Speech Processing, robotics etc. due to their high accuracies. But to use these neural networks on embedded device is a challenging task as the models are very large in terms of memory footprint and computational complexity. This work proposes two approaches: i) combining pruning which removes the unimportant weights with structured sparsity learning which trims the structure of neural network to preserve the density of computations, and ii) a novel approach on combining quantized neural network i.e. networks with weights of smaller bit-width with pruning i.e. removing the unimportant weights from the network. The experiments showed results on large networks such as VGG inspired ConvNet on SVHN dataset compressed to 79x for a 12% loss in the accuracy.

3.3.2.15 Parallelization

The work employed data parallel neural network training.

3.3.2.16 SERC Resources and Experiments

Experiment were run on one NVIDIA DGX-1 node of SERC.

3.3.2.17 Performance and Scale

We used department cluster for small-scale neural networks. Large-scale neural networks could only be trained in the NVIDIA DGX-1 machine of SERC.

3.3.2.18 Research

Initial Explorations

Manasi Twari, Renga Bashyam, Sathish Vadhiyar

Initial explorations were also performed in the following areas:

1. K-Nearest Neighbor (k-NN) search is one of the most commonly used approaches for similarity search. It finds extensive applications in machine learning and data mining. In this paper, we propose a solution towards this end where we use vantage point trees for partitioning the dataset across multiple processes and exploit an existing graph-based sequential approximate k-NN search algorithm called HNSW (Hierarchical Navigable Small World) for searching locally within a process. Executions were performed on SahasraT CPU up to 8192 cores.
2. A novel pipelined Conjugate Gradient solver was developed. Executions were performed on SahasraT CPU up to 1600 cores.

3.3.2.19 Students Graduated

1. Mayank Gupta, MTech (CDS)
2. Udit Gupta, MTech (CDS)

3.3.3 Prof. Venkatesh Babu's Lab

3.3.3.1 Research

The lab used Nvidia DGX machine for training Convolutional Neural Networks (CNN) for following different tasks:

1. Training a CNN model to perform HDR image fusion given an input sequence of varying exposure images.
2. Training a few-shot domain adaptation CNN model for inverse imaging problems.
3. Training robust Deep Neural Networks.
4. Training CNNs to detect persons in densely crowded images. This involves training complex models that require high computational power and memory. To report metrics over large datasets, complex network training takes a few days to run. Thus, the SERC DGX GPU cluster with parallelization and larger GPU memory has enabled this training and cut short training time to hours.

3.3.3.2 SERC Resources and Experiments

Nvidia DGX machine: Jobs used 1,2 and 4 GPUs in parallel with up to 16 nodes with 8 CPU cores per node.

3.3.3.3 Parallelization

- a. Idea is to perform smoothing iterations of a certain type of DOFs on GPU at the same time perform smoothing iterations for other type of DOFs on host CPU in order to utilize both GPU and CPU resources simultaneously by leveraging upon the asynchronous nature of CUDA API calls and CUDA streams.
- b. OpenMP for loop parallelization is also employed in the smoothing iteration step and DOF coloring step so that DOFs of the same color can be processed in parallel.
- c. CUDA Multi-Process service (MPS) is used to enable concurrent execution of kernel and memcopy operations from different MPI processes on the same GPU device on a node.

3.3.3.4 Performance

There is a 1.5-2x speedup when DGX is used. For training on the ImageNet dataset, it is not possible to use a single GPU, whereas on DGX we could run it across 4 GPUs in <48hrs. For typical runs, a 12hr job would finish in 7-8hrs. Our lab GPUs cannot be used for running large-size CNN models. These experiments could only be run on the Nvidia DGX machine

3.3.3.5 Outcomes



- a. Towards Achieving Adversarial Robustness by Enforcing Feature Consistency Across Bit Planes, CVPR 2020
- b. Locate, Size and Count: Accurately Resolving People in Dense Crowds via Detection, TPAMI 2020.

3.4 Department of Chemical Engineering (CE)


3.4.1 Prof. Prabhu Nott's Lab

3.4.1.1 Research

The flow of granular materials, such as food grains, pharmaceutical powders, mineral ores and construction materials, are of immense importance to industry, but a fundamental understanding of the process is lacking. To build such an understanding, we utilize particle dynamics simulations, which have helped us validate mathematical models, and understand flow in devices such as screw feeders that are widely used in the industry.



DEM Simulations of simple shear: Understanding flow behavior in simple systems



DEM Simulation of a screw feeder: SERC's HPC systems help us in probing flows in more complicated systems.

3.4.1.2 SERC Resources and Experiments

The Sahasrat HPC is used with core counts ranging from 240 to 1600.

3.4.1.3 Parallelization

We implement our simulations using LAMMPS (a module on the Sahasrat machine) that uses MPI to parallelize the code.

3.4.1.4 Performance

On the HPC: a few hours. On the desktop PC: Either a few days or unviable.

3.4.1.5 Outcomes

Publications:

Dsouza, P., & Nott, P. (2020). A non-local constitutive model for slow granular flow that incorporates dilatancy. *Journal of Fluid Mechanics*, 888, R3. doi:10.1017/jfm.2020.62

Industry Impact

Executed a project funded by Unilever with partial use of Saharat. Currently operating projects funded by Boeing and the International Fine Particle Research Institute where usage of SERC HPC is envisaged.

Student Graduations:

Aashish Gupta (MSc), Krishnaraj KP (PhD)

3.5 Centre for Earth Sciences (CEaS)

3.5.1 Dr Attreyee Ghosh's Lab

3.5.1.1 Research

The lab is investigating the causes behind the survival of cratons, the oldest parts of the planet. These regions have survived for billions of years as opposed to other parts of the Earth that gets destroyed within a few million years. The reasons for their survival are still unknown and in this project, with the help of numerical modeling, the lab is exploring some of the physical parameters that might have enabled these cratons to survive for such a long time. The deformation of the Indian plate is also being investigated using numerical models

3.5.1.2 SERC Resources and Experiments

Sahasrat, maximum cores used: 768.

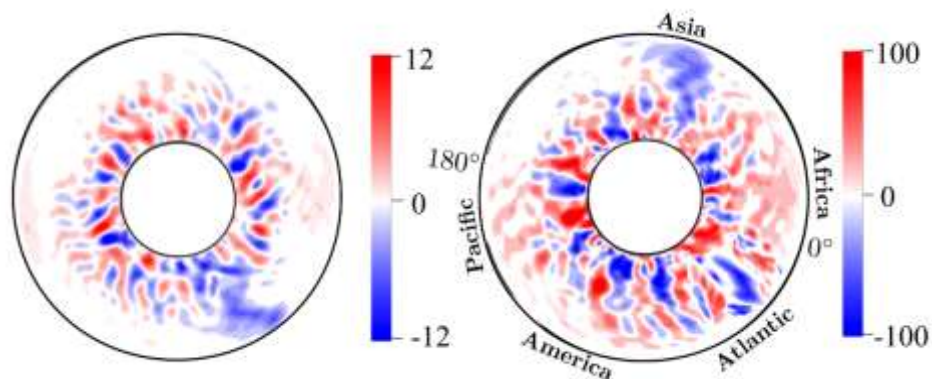
3.5.1.3 Outcomes

- J. Paul, A. Ghosh, Evolution of cratons through the ages: A time-dependent study, *Earth & Planetary Science Letters*, 531, <https://doi.org/10.1016/j.epsl.2019.115962>, 2020.
- J. Paul*, A. Ghosh, C. P. Conrad, Traction and strain-rate at the base of the lithosphere: An insight into cratonic survival, *Geophysical Journal International*, 217, p. 1024-1033, 2019.
- Video https://www.youtube.com/playlist?list=UUBruY_7eCEfVjnlriZJGdqw.

3.5.2 Prof. Binod Sreenivasan's Lab

3.5.2.1 Research

The magnetic fields of Earth and planets are generated by dynamo action in their fluid cores. The numerical simulation of the planetary dynamo problem involves the solution of the nonlinear magnetohydrodynamic (MHD) equations on massively parallel computers. The problems solved by us on SERC's supercomputer SahasraT in 2019 are devoted to understanding (a) the preference for the axial dipole in rotating planetary dynamos, and (b) the effect of large lower-mantle heterogeneity on the geodynamo.



Contours of the radial velocity at the equator in a spherical shell simulation of convection in the Earth's core subject to a large lower mantle heterogeneity. A coherent downwelling that forms beneath the Atlantic in a mildly driven convection (left panel) switches over to Asia in strongly driven convection (right panel). These simulations may explain the eastward migration of the magnetic North Pole, observed in the present-day geomagnetic field.

3.5.2.2 *SERC Resources and Experiments*

SahasraT, typically 2160 cores.

3.5.2.3 *Parallelization*

The computations are performed in a spherical shell geometry, with MPI parallelization in the radial and latitudinal grid points.

3.5.2.4 *Performance*

One simulation involves 8-10 submissions on a 24-hour queue in SahasraT. The simulations are up to 8 times faster than those on the AMD 6378 cluster on 256 CPUs.

3.5.2.5 *Outcomes*

- a. S. Sahoo and B. Sreenivasan, Convection in a rapidly rotating cylindrical annulus with laterally varying boundary heat flux, *Journal of Fluid Mechanics*, v883, A1, 2019.
doi:10.1017/jfm.2019.803.
- b. Indo-French Centre for the Promotion of Advanced Research, Grant No. 5307-1.

3.6 Department of Electronic Systems Engineering (DESE)

3.6.1 Prof. T.V. Prabhakar's Lab

3.6.1.1 *Research*

Speech is most prominent and one of the natural forms of communication among of human being. Sometimes reading becomes more powerful than the listening ability, and speech-to-text conversion fits as an important tool in such cases. The purpose of this project is to study and develop a speech recognition system for a cabin environment to decode the announcements made during the flight. We develop speech recognition systems using open source platforms like TensorFlow. We develop neural network based solutions using TensorFlow.

3.6.1.2 *SERC Resources and Experiments*

Up to 4 GPU cores of the NVIDIA-DGX GPU cluster.

3.6.1.3 *Parallelization*

GPU version of the TensorFlow platform.

3.6.1.4 *Performance*

The execution on normal CPU based systems would run for around a month and a half, but it would complete within a week on the NVIDIA-DGX machine.

3.7 Inorganic and Physical Chemistry (IPC)


3.7.1 Prof. E.D. Jemmis' Lab

3.7.1.1 Research

We study the structure and reactivity problems of real-life molecules, clusters and solids using theoretical techniques ranging from the simplest of molecular orbital methods to the sophisticated electronic structure theory depending on the system at hand and the questions that are to be answered. Special emphasis is placed in weaving threads between problems in one area to another; between polymorphs of elements and their compounds, between organic and organometallic chemistry, amongst the chemistry of various main group elements; Bonding, Structure and Reactions across the Periodic Table of Elements. We place great importance in not only getting numbers as an answer to a problem, but also in finding out why the numbers turn out the way they do, based on overlap of orbitals, perturbation theory, and symmetry, and in devising transferable models.

Current emphasis of Research is in the following areas:

- Chemistry of elemental boron, boron clusters, boranes, metallaboranes and metal borides
- Transition metal organometallics, activation of small molecules, C-H bond activation, agostic interactions, C-C coupling reactions.
- H-bond, and weak interactions across the periodic table
- Chemistry of heavier main group elements



The stability of di-dehydrogeno polyhedral boranes is determined by the overlap of exohedral orbitals, which, in turn, depends on the size of polyhedral.

3.7.1.2 SERC Resources and Experiments

Solid state calculations have been done using VASP, installed in SAHASRAT. The number of node count varies from 5 to 20. Parallelization strategy is used to enhance the performance of the code. Most of the jobs generally take two to three days to complete, as the number of atoms in most of the cases is more than 100.

1.85 million core hours of SahasraT were used.

The lab has used Gaussian 09 program package installed in dell, delta and tyrone clusters for molecular electronic structure calculation. The structural optimization and electronic structure calculations are done to understand the stability of different structures. The intermediate and transition state calculations are also performed. The number of nodes used for the calculations were 1-4 and number of processors used were 16-256 as given in SERC site.

Total CPU time (in CPU hours): Tyrone cluster-11960059, Delta cluster-7555023, Dell-111080206.

3.7.1.3 Performance

Particularly, tyrone cluster with 256 core and delta cluster with 32 core are found to be quite faster than the IPC-cluster with 8 processors (max. limit).

3.7.1.4 Publications and outcomes

1. Designing M-bond ($X-M \cdots Y$, M = transition metal): σ -hole and radial density distribution. J. Joy, E.D. Jemmis. *J. Chem. Sci.* 131 (2019) 1-8.
2. Organoaluminum cations for carbonyl activation. R. Kannan, R. Chambenahalli, S. Kumar, A. Krishna, A.P. Andrews, E.D. Jemmis, A. Venugopal. *Chem. Commun.* 55 (2019) 14629-14632.
3. Stabilization of Classical $[B_2H_5]^-$: Structure and Bonding of $[(Cp^*Ta)_2(B_2H_5)(\mu-H)L_2]$ ($Cp^* = \eta^5-C_5Me_5$; $L = SCH_2S$). K. Saha, S. Ghorai, S. Kar, S. Saha, R. Halder, B. Raghavendra, E. D. Jemmis, S. Ghosh. *Angew. Chem. Int. Ed.* 58 (2019) 17684-17689.
4. Isolation of base stabilized fluoroborylene and its radical cation. S.K. Sarkar, M.M. Siddiqui, S. Kundu, M. Ghosh, J. Kretsch, P. Stollberg, R. Herbst-Irmer, D. Stalke, A.C. Stückl, B. Schwederski, W. Kaim, S. Ghorai, E.D. Jemmis, H.W. Roesky. *Dalt. Trans.* 48 (2019) 8551-8555.
5. A Dicationic Bismuth(III) Lewis Acid: Catalytic Hydrosilylation of Olefins. S. Balasubramaniam, S. Kumar, A.P. Andrews, B. Varghese, E.D. Jemmis, A. Venugopal. *Eur. J. Inorg. Chem.* 2019 (2019) 3265-3269.
6. A theoretical analysis of the structure and properties of B_26H_{30} isomers: Consequences to the laser and semiconductor doping capabilities of large borane clusters. Macháček, Jan, Antonio Francés-Monerris, Naiwrit Karmodak, Daniel Roca-Sanjuán, Jindřich Fanfrlík, Michael GS Londesborough, Drahommír Hnyk, and Eluvathingal D. Jemmis. *Phys. Chem. Chem. Phys.* 21 (2019) 12916-12923.
7. Overlap of Radial Dangling Orbitals Controls the Relative Stabilities of Polyhedral B_nH_n-x Isomers ($n = 5-12$, $x = 0$ to $n - 1$). N. Karmodak, R. Chaliha, E.D. Jemmis. *Inorg. Chem.* 58 (2019) 3627-3634.

Manpower trained: One PhD student, Dr. Naiwrit Karmodak, completed PhD during this period. His thesis was adjudged to be the best thesis of the Dept for the year 2019. Two int-PhD students and two regular PhD

students continue their research using SERC facilities. Three postdoctoral fellows also use the facility for research.

One IISc-UG student, Arindam Sau completed BSc final year project and is about to submit MS thesis. Facilities at SERC helped enormously in his work. Though to a much less degree, several short term summer research participants get a flavor of computational chemistry research using the facilities at IPC and on occasions from the SERC facilities.

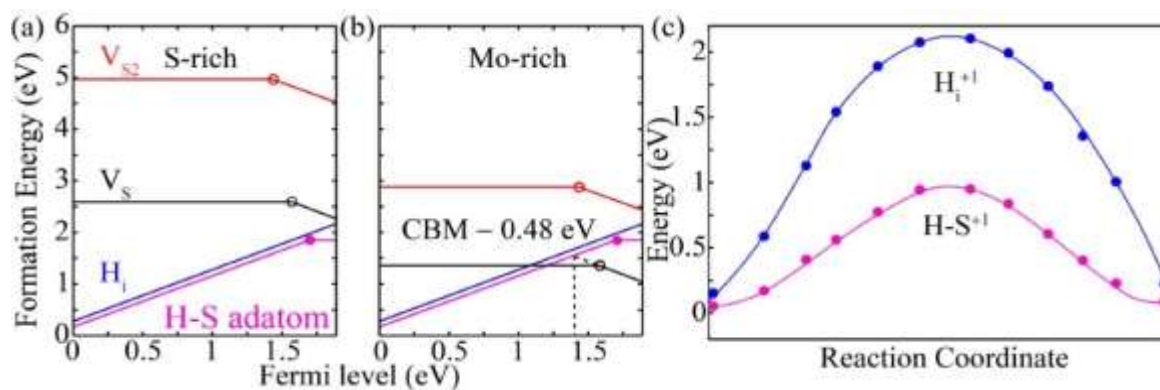

3.8 Materials Research Centre (MRC)

3.8.1 Prof. Abhishek Singh's Lab


3.8.1.1 Research

The lab has used SERC HPC facility for computational modelling of materials, especially in applications, where extensive computational resources were required. The problems solved using this facility include understanding of material behavior at atomistic scale and determining its physical and chemical properties using first principles density functional theory (DFT). These include understanding the origin of defects in materials, evaluation of various contributions to the lattice thermal conductivity of thermoelectric materials and thermoelectric figure of merit, discovery of novel topological materials with unique properties, and providing the insights into catalytic activities involved in some of the key reactions. The HPC system also provides a great support to generate several dataset required for high throughput screening of materials with desired property as well as development of machine learning prediction models.

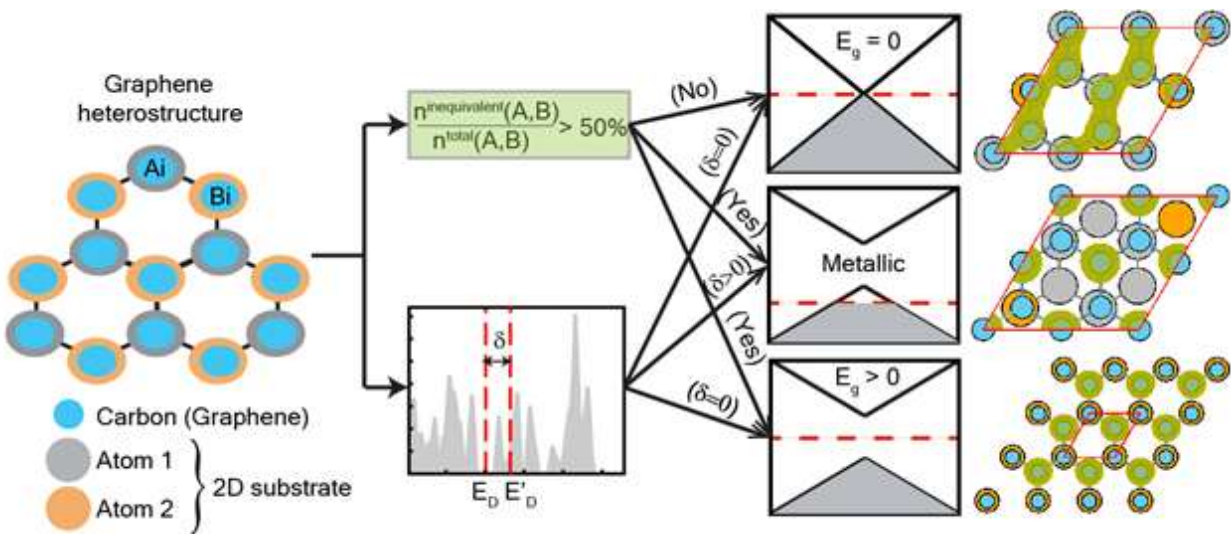
Origin of n-type conductivity of monolayer MoS₂



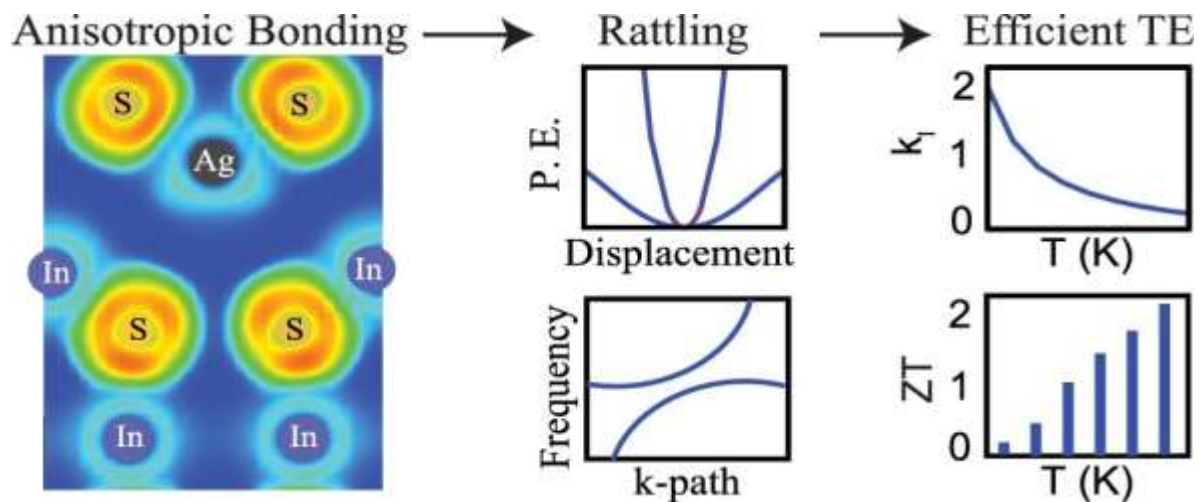
Topological Phases in Hydrogenated Group 13 Monolayers



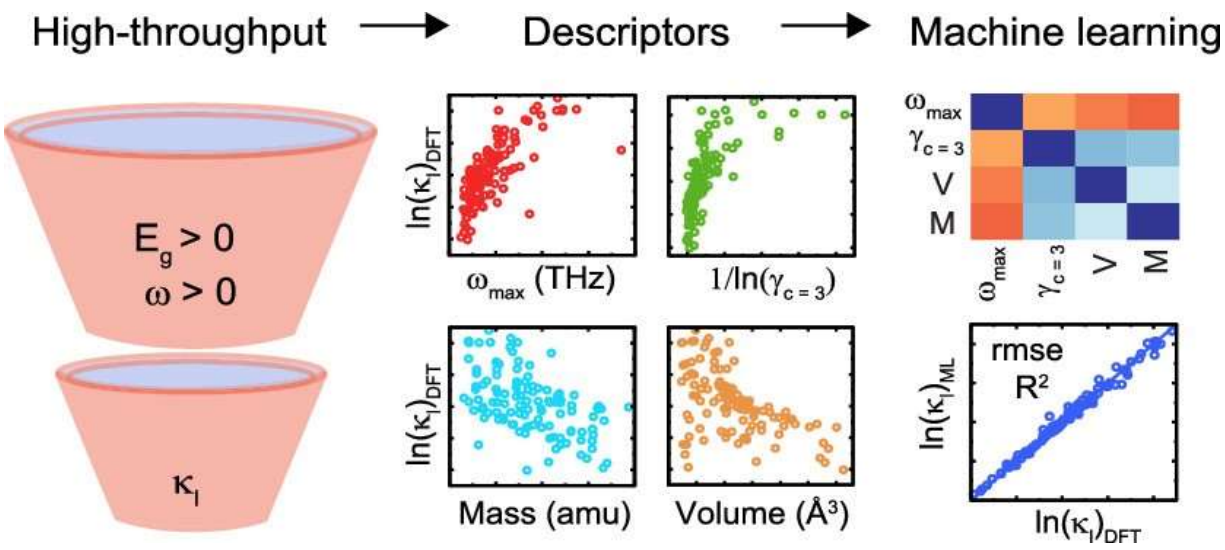
Critical Sublattice Symmetry Breaking



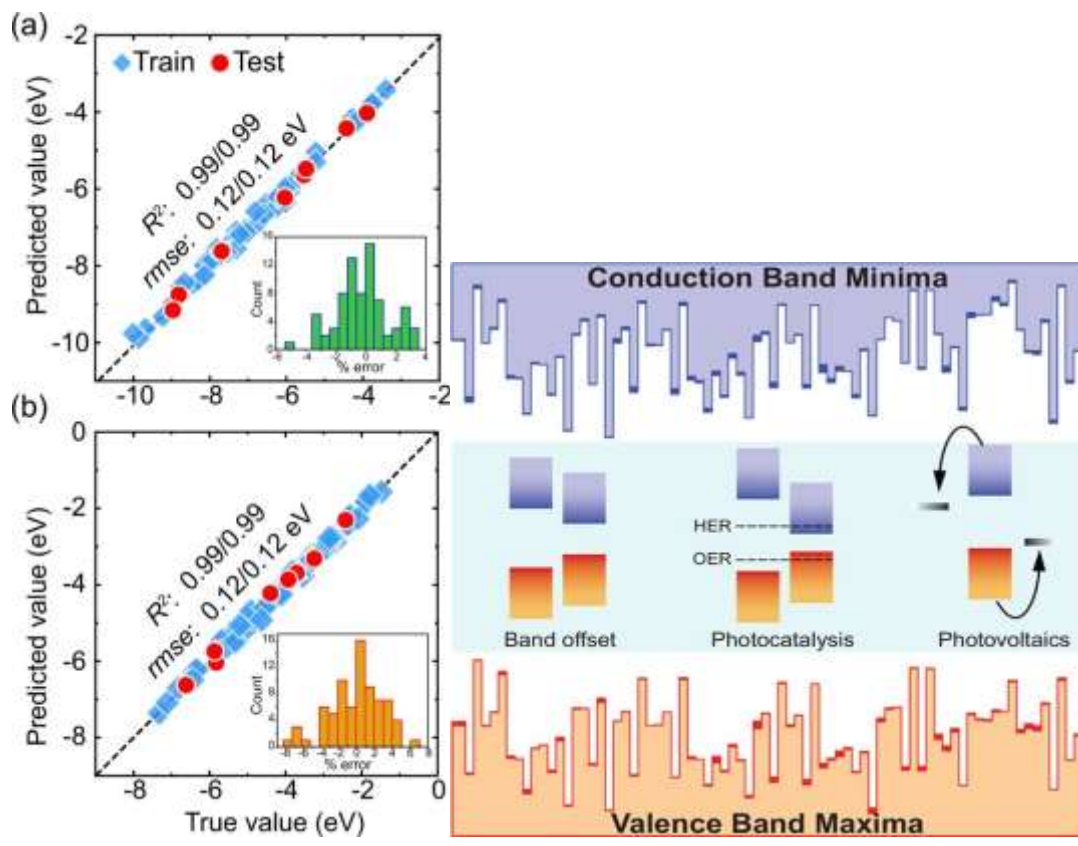
Rattling-Induced Ultra-low Thermal Conductivity



High-throughput Property Map to Machine Learning for Predicting Lattice Thermal Conductivity



Accelerated Data-driven Accurate Positioning of the Band-edges of MXenes



3.8.1.2 SERC Resources and Experiments

In most of the problems, the commonly used HPC resources are debug, small and small72 queue with 240 cores. However, in few cases, we have also utilized medium queue.

A parallel code VASP, Quantum Espresso, WannierTools, ShengBTE were used, with both multicore and multinode parallelization.

3.8.1.3 Performance

The computational facilities at SERC are much faster than our lab cluster. Typical execution times using SERC HPC systems were 24 hours or 72 hours. The same task would have taken at least an order of magnitude more time on the local facility.

3.8.1.4 Publications

1. Noninvasive Subsurface Electrical Probe for Encapsulated Layers in van der Waals Heterostructures, M. Pandey, R. Soni, A. Mathur, A. Singh, A. K. Singh, S. Raghavan, and U. Chandni, *Phys. Rev. A*, **12**, 064032, (2019).
2. Revealing carbon mediated luminescence centers with enhanced lifetime in porous alumina, S. Bhowmick, S. Pal, A. Singh, M. Gupta, D. M. Phase, A. K. Singh, A. Kanjilal, *J. Appl. Phys.* **126**, 164904, (2019).
3. Topological Phases in Hydrogenated Group 13 Monolayers, R. K. Barik, R. Kumar, A. K. Singh, *J. Phys. Chem. C*, **123**, 42, 25985-25990, (2019).
4. Critical Sublattice Symmetry Breaking: A Universal Criterion for Dirac Cone Splitting, R. Kumar, D. Das, E. Munoz, A. K. Singh, *J. Phys. Chem. C*, **123**, 37, 23082-23088, (2019).
5. Rattling-Induced Ultra-low Thermal Conductivity Leading to Exceptional Thermoelectric Performance in AgIn5S8, R. Juneja, A. K. Singh, *ACS Appl. Mater. & Interfaces*, **11**, 37, 33894-33900, (2019).
6. Coupling High-throughput Property Map to Machine Learning for Predicting Lattice Thermal Conductivity, R. Juneja, G. Yumnam, S. Satsangi, A. K. Singh, *Chem. Mater.*, **31**, 14, 5145-5151, (2019).
7. Structure-dependent electrical and magnetic properties of iron oxide composites, N. Lertcumfu, F. N. Sayed, S. Shirodkar, S. Radhakrishnana, A. Mishra, G. Rujijanagul, A. K. Singh, B. I. Yakobson, C. S. Tiwary, P. M. Ajayan, *physica status solidi (a)*, **216**, 16, 1801004, (2019).
8. Magnetism in two dimensional materials beyond Graphene, N. Sethulakshmi A. Mishra, P M Ajayan, Y. Kawazoe, A. K. Roy, A. K. Singh, C. S. Tiwary, *Mater. Today*, **27**, 107-122, (2019).
9. Thermal Conductivity Enhancement in MoS₂ under Extreme Strain, X. Meng, T. Pandey, J. Jeong, S. Fu, J. Yang, K. Chen, A. Singh, F. He, X. Xu, J. Zhou, W-P Hsieh, A. K. Singh, J-F Lin, Y. Wang, *Phys. Rev. Lett.*, **122**, 155901, (2019).
10. Origin of Ultralow Thermal Conductivity in n-Type Cubic Bulk AgBiS₂: Soft Ag Vibrations and Local Structural Distortion Induced by the Bi 6s² Lone Pair, E. Rathore, R. Juneja, S. P. Culver, N. Minafra, A. K. Singh, W. G. Zeier, and K. Biswas, *Chem. Mater.* **31**, 6, 2106-2113, (2019).
11. Recent advances in MXenes: From fundamentals to applications, M. Khazaei, A. Mishra, N. S. Venkataramanan, A. K. Singh and S. Yunoki, *Current Opinion in Solid State & Materials Science*, **23**, **3**, 164-178, (2019).
12. Origin of n-type conductivity of monolayer MoS₂, A. Singh and A. K. Singh, *Phys. Rev. B: Rapid Comm.*, **99**, 121201(R), (2019).
13. Accelerated Data-driven Accurate Positioning of the Band-edges of MXenes, A. Mishra, S. Satsangi, A. C. Rajan, H. Mizuseki, K. R. Lee and A. K. Singh, *J. Phys. Chem. Lett.*, **10**, 780, (2019).
14. Morphology controlled synthesis of low bandgap SnSe₂ with high photodetectivity, R. K. Rai, S. Islam, A. Roy, G. Agrawal, A. K. Singh, A. Ghosh and N. Ravishankar, *Nanoscale*, **11**, 870, (2019).

15. Amine Functionalized Zirconium Metal-Organic Framework as an Effective Chemiresistive Sensor for Acidic Gases, M. E. DMello, N. Sundaram, A. Singh, A. K. Singh and S. B. Kalidindi, *Chem. Commun.* **55**, 349, (2019)

3.9 Department of Physics (PHY)

3.9.1 Prof. Rahul Pandit's Lab

3.9.1.1 Research

Lab members have studied a wide variety of problems over the last year. These are in the areas of (a) turbulence in nonlinear hydrodynamical partial differential equations, (b) spiral- and scroll-wave turbulence in mathematical models for cardiac tissue, and (c) localization in quasiperiodic systems.

3.9.1.2 Publications

1. A first-passage-time problem for tracers in homogeneous and isotropic fluid turbulence
AK Verma, A Bhatnagar, D Mitra, R Pandit
arXiv preprint arXiv:2001.01260
2. The formation of compact objects at finite temperatures in a dark-matter-candidate self-gravitating bosonic system
AK Verma, R Pandit, ME Brachet
arXiv preprint arXiv:1912.10172
3. Path-planning microswimmers can swim efficiently in turbulent flows
JK Alageshan, AK Verma, J Bec, R Pandit
arXiv preprint arXiv:1910.01728
4. Detection and Termination of Broken-Spiral-Waves in Mathematical Models for Cardiac Tissue: A Deep-Learning Approach
MK Mulimani, JK Alageshan, R Pandit
2019 Computing in Cardiology (CinC), Page 1-Page 4
5. The one-dimensional Kardar-Parisi-Zhang and Kuramoto-Sivashinsky universality class: limit distributions
D Roy, R Pandit
arXiv preprint arXiv:1908.06007
6. Two-dimensional magnetohydrodynamic turbulence with large and small energy-injection length scales
D Banerjee, R Pandit
Physics of Fluids 31 (6), 065111

7. Transport, multifractality, and the breakdown of single-parameter scaling at the localization transition in quasiperiodic systems
J Sutradhar, S Mukerjee, R Pandit, S Banerjee
Physical Review B 99 (22), 224204
8. Deep-learning-assisted detection and termination of spiral-and broken-spiral waves in mathematical models for cardiac tissue
MK Mulimani, JK Alageshan, R Pandit
arXiv preprint arXiv:1905.06547
9. The Statistical Properties of Superfluid Turbulence in He from the Hall-Vinen-Bekharevich-Khalatnikov Model
AK Verma, V Shukla, A Basu, R Pandit
arXiv preprint arXiv:1905.01507
10. Anisotropic shortening in the wavelength of electrical waves promotes onset of electrical turbulence in cardiac tissue: an *in silico* study, S. Zimik, and R. Majumder, accepted for publication in PLOS ONE (2020).

4. Access to External Users

In September 2019, SERC evolved a policy to provide access, at a charge, for Researchers and Engineers from Academic organizations, Government-funded R&D laboratories and Industry to state-of-the-art compute facilities such as SahasraT, DGX-1 and future high-end procurements for R&D work. Towards the end of 2019, 3 proposals were approved from faculties at IITs and Government Colleges. Resource is slated to be utilised in 2020.

5. Publications in 2019 using SahasraT

1. Rohit Singhal, S. Ravichandran, Sourabh S. Diwan and Garry L. Brown, *Reynolds stress gradient and vorticity fluxes in axisymmetric turbulent jet and plume*, The 16th Asian Congress of Fluid Mechanics, 13-17 December 2019, JNCASR Jakkur Bengaluru, India.
2. J. Paul*, A. Ghosh, C. P. Conrad, Traction and strain-rate at the base of the lithosphere: An insight into cratonic survival, *Geophysical Journal International*, 217, p. 1024-1033, 2019.
3. S. Sahoo and B. Sreenivasan, Convection in a rapidly rotating cylindrical annulus with laterally varying boundary heat flux, *Journal of Fluid Mechanics*, v883, A1, 2019. doi:10.1017/jfm.2019.803.
4. Designing M-bond (X-M··· Y, M = transition metal): σ -hole and radial density distribution. J. Joy, E.D. Jemmis. *J. Chem. Sci.* 131 (2019) 1-8.
5. Organoaluminum cations for carbonyl activation. R. Kannan, R. Chambenahalli, S. Kumar, A. Krishna, A.P. Andrews, E.D. Jemmis, A. Venugopal. *Chem. Commun.* 55 (2019) 14629-14632.
6. Stabilization of Classical [B₂H₅]⁻: Structure and Bonding of [(Cp*Ta)₂(B₂H₅)(μ -H)L₂] (Cp*= η ⁵-C₅Me₅; L=Sch2S). K. Saha, S. Ghorai, S. Kar, S. Saha, R. Halder, B. Raghavendra, E. D. Jemmis, S. Ghosh. *Angew. Chem. Int. Ed.* 58 (2019) 17684-17689.

7. Isolation of base stabilized fluoroborylene and its radical cation. S.K. Sarkar, M.M. Siddiqui, S. Kundu, M. Ghosh, J. Kretsch, P. Stollberg, R. Herbst-Irmer, D. Stalke, A.C. Stückl, B. Schwederski, W. Kaim, S. Ghorai, E.D. Jemmis, H.W. Roesky. *Dalt. Trans.* 48 (2019) 8551-8555.
8. A Dicationic Bismuth(III) Lewis Acid: Catalytic Hydrosilylation of Olefins. S. Balasubramaniam, S. Kumar, A.P. Andrews, B. Varghese, E.D. Jemmis, A. Venugopal. *Eur. J. Inorg. Chem.* 2019 (2019) 3265-3269.
9. A theoretical analysis of the structure and properties of B₂₆H₃₀ isomers: Consequences to the laser and semiconductor doping capabilities of large borane clusters. Macháček, Jan, Antonio Francés-Monerris, Naiwrit Karmodak, Daniel Roca-Sanjuán, Jindřich Fanfrlík, Michael GS Londesborough, Drahommír Hnyk, and Eluvathingal D. Jemmis. *Phys. Chem. Chem. Phys.* 21 (2019) 12916-12923.
10. Overlap of Radial Dangling Orbitals Controls the Relative Stabilities of Polyhedral B_nH_{n-x} Isomers (n= 5-12, x = 0 to n - 1). N. Karmodak, R. Chaliha, E.D. Jemmis. *Inorg. Chem.* 58 (2019) 3627-3634.
11. Noninvasive Subsurface Electrical Probe for Encapsulated Layers in van der Waals Heterostructures, M. Pandey, R. Soni, A. Mathur, A. Singh, A. K. Singh, S. Raghavan, and U. Chandni, *Phys. Rev. A*, **12**, 064032, (2019).
12. Revealing carbon mediated luminescence centers with enhanced lifetime in porous alumina, S. Bhowmick, S. Pal, A. Singh, M. Gupta, D. M. Phase, A. K. Singh, A. Kanjilal, *J. Appl. Phys.* **126**, 164904, (2019).
13. Topological Phases in Hydrogenated Group 13 Monolayers, R. K. Barik, R. Kumar, A. K. Singh, *J. Phys. Chem. C*, **123**, 42, 25985-25990, (2019).
14. Critical Sublattice Symmetry Breaking: A Universal Criterion for Dirac Cone Splitting, R. Kumar, D. Das, E. Munoz, A. K. Singh, *J. Phys. Chem. C*, **123**, 37, 23082-23088, (2019).
15. Rattling-Induced Ultra-low Thermal Conductivity Leading to Exceptional Thermoelectric Performance in AgIn₅S₈, R. Juneja, A. K. Singh, *ACS Appl. Mater. & Interfaces*, **11**, 37, 33894-33900, (2019).
16. Coupling High-throughput Property Map to Machine Learning for Predicting Lattice Thermal Conductivity, R. Juneja, G. Yumnam, S. Satsangi, A. K. Singh, *Chem. Mater.*, **31**, 14, 5145-5151, (2019).
17. Structure-dependent electrical and magnetic properties of iron oxide composites, N. Lertcumfu, F. N. Sayed, S. Shirodkar, S. Radhakrishnana, A. Mishra, G. Rujijanagul, A. K. Singh, B. I. Yakobson, C. S. Tiwary, P. M. Ajayan, *physica status solidi (a)*, **216**, 16, 1801004, (2019).
18. Magnetism in two dimensional materials beyond Graphene, N. Sethulakshmi A. Mishra, P M Ajayan, Y. Kawazoe, A. K. Roy, A. K. Singh, C. S. Tiwary, *Mater. Today*, **27**, 107-122, (2019).
19. Thermal Conductivity Enhancement in MoS₂ under Extreme Strain, X. Meng, T. Pandey, J. Jeong, S. Fu, J. Yang, K. Chen, A. Singh, F. He, X. Xu, J. Zhou, W-P Hsieh, A. K. Singh, J-F Lin, Y. Wang, *Phys. Rev. Lett.*, **122**, 155901, (2019).
20. Origin of Ultralow Thermal Conductivity in n-Type Cubic Bulk AgBiS₂: Soft Ag Vibrations and Local Structural Distortion Induced by the Bi 6s₂ Lone Pair, E. Rathore, R. Juneja, S. P. Culver, N. Minafra, A. K. Singh, W. G. Zeier, and K. Biswas, *Chem. Mater.* **31**, 6, 2106-2113, (2019).
21. Recent advances in MXenes: From fundamentals to applications, M. Khazaei, A. Mishra, N. S. Venkataramanan, A. K. Singh and S. Yunoki, *Current Opinion in Solid State & Materials Science*, **23**, **3**, 164-178, (2019).
22. Origin of n-type conductivity of monolayer MoS₂, A. Singh and A. K. Singh, *Phys. Rev. B: Rapid Comm.*, **99**, 121201(R), (2019).

23. Accelerated Data-driven Accurate Positioning of the Band-edges of MXenes, A. Mishra, S. Satsangi, A. C. Rajan, H. Mizuseki, K. R. Lee and A. K. Singh, *J. Phys. Chem. Lett.*, **10**, 780, (2019).
24. Morphology controlled synthesis of low bandgap SnSe₂ with high photodetectivity, R. K. Rai, S. Islam, A. Roy, G. Agrawal, A. K. Singh, A. Ghosh and N. Ravishankar, *Nanoscale*, **11**, 870, (2019).
25. Amine Functionalized Zirconium Metal-Organic Framework as an Effective Chemiresistive Sensor for Acidic Gases, M. E. DMello, N. Sundaram, A. Singh, A. K. Singh and S. B. Kalidindi, *Chem. Commun.* **55**, 349, (2019)
26. A first-passage-time problem for tracers in homogeneous and isotropic fluid turbulence
AK Verma, A Bhatnagar, D Mitra, R Pandit
arXiv preprint arXiv:2001.01260
27. The formation of compact objects at finite temperatures in a dark-matter-candidate self-gravitating bosonic system
AK Verma, R Pandit, ME Brachet
arXiv preprint arXiv:1912.10172
28. Path-planning microswimmers can swim efficiently in turbulent flows
JK Alageshan, AK Verma, J Bec, R Pandit
arXiv preprint arXiv:1910.01728
29. Detection and Termination of Broken-Spiral-Waves in Mathematical Models for Cardiac Tissue: A Deep-Learning Approach
MK Mulimani, JK Alageshan, R Pandit
2019 Computing in Cardiology (CinC), Page 1-Page 4
30. The one-dimensional Kardar-Parisi-Zhang and Kuramoto-Sivashinsky universality class: limit distributions
D Roy, R Pandit
arXiv preprint arXiv:1908.06007
31. Two-dimensional magnetohydrodynamic turbulence with large and small energy-injection length scales
D Banerjee, R Pandit
Physics of Fluids 31 (6), 065111
32. Transport, multifractality, and the breakdown of single-parameter scaling at the localization transition in quasiperiodic systems
J Sutradhar, S Mukerjee, R Pandit, S Banerjee
Physical Review B 99 (22), 224204

33. Deep-learning-assisted detection and termination of spiral-and broken-spiral waves in mathematical models for cardiac tissue
MK Mulimani, JK Alageshan, R Pandit
arXiv preprint arXiv:1905.06547
34. The Statistical Properties of Superfluid Turbulence in He from the Hall-Vinen-Bekharevich-Khalatnikov Model
AK Verma, V Shukla, A Basu, R Pandit
arXiv preprint arXiv:1905.01507
35. Rintu Panja, Sathish S. Vadhiyar: HyPar: A divide-and-conquer model for hybrid CPU-GPU graph processing. *J. Parallel Distributed Comput.* 132: 8-20 (2019).
36. Anwesa Bhowmik, Sathish Vadhiyar: HyDetect: A Hybrid CPU-GPU Algorithm for Community Detection. [HiPC 2019](#): 2-11.
37. Udit Gupta, Sathish Vadhiyar: Fast and Accurate Learning of Knowledge Graph Embeddings at Scale. *HiPC 2019*: 173-182.

6. Projects Based on SahasraT

1. Indo-French Centre for the Promotion of Advanced Research, Grant No. 5307-1.
2. Project funded by Unilever with partial use of Saharat