HPC System software, Management, benchmarking and Application Software Initiatives in India under National Supercomputing Mission

Sanjay Wandhekar, Senior Director, Head HPC Technologies
C-DAC

About C-DAC

Established in 1988

Presence at 12 locations

Key Areas

- HPC & Quantum Computing
- DIR-V and Strategic Electronics
- Software Technologies including FOSS
- Health Informatics
- Cyber Security & Cyber Forensics
- Multilingual & Heritage Computing
- Education and Training

Build indigenous supercomputing capability

National Supercomputing Mission (NSM)

One Vision One Goal…Advanced Computing for Human advancement…

National Supercomputing Mission

Creation of Supercomputing Infrastructure

- **15 Supercomputers with compute capacity of 24 PF established**
- **9 more Supercomputers with compute capacity of 40 PF to be build in next 12-15 months – Includes a 20PF National facility**

Supercomputing Infrastructure usage till date:

- **6700+ HPC users from 120+ institutes across the country**
- **84+ Lakhs HPC jobs executed**
- **500+ PhD students**
- **850+ publications**

Development of Indigenous Technologies

Compute node, HPC N/w and Software stack

HPC applications for National Need in 5 domains

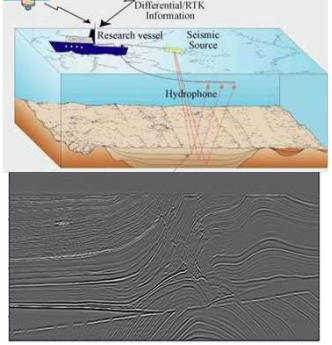

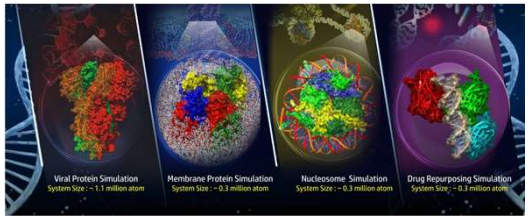

HPC Human Resource Development

- **19500+ HPC Human Resources trained**

A Glimpse of NSM Systems

AIRAWAT –PARAM SIDDHI-AI Largest Supercomputer in India

- **AIRAWAT PSAI of 8.5 Petaflops (410 AI Petaflops, Peak 13.17 Petaflops) is the fastest Supercomputer in India and ranked at No. 75 position in 'TOP500 Supercomputer List – June 2023' declared at ISC 23 at Germany.**
- **600+ Users from 100+ institutes**

NSM Indigenous Developments

- Rudra series of Servers
- HPC Processor AUM
- Trinetra Scalable HPC Network
- System Software Stack – HPC & AI
- Storage – DAOS based Flash Storage
- Liquid Cooling technology
- Experience in architecting and building 10’s of Petaflop Supercomputer with efficient Data centre infrastructure

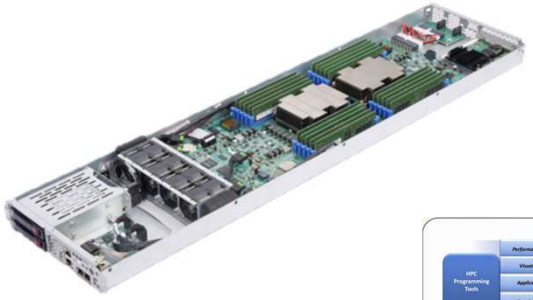

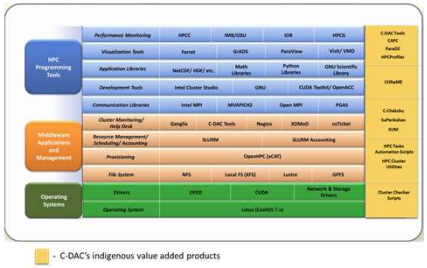

HPC Software STACK

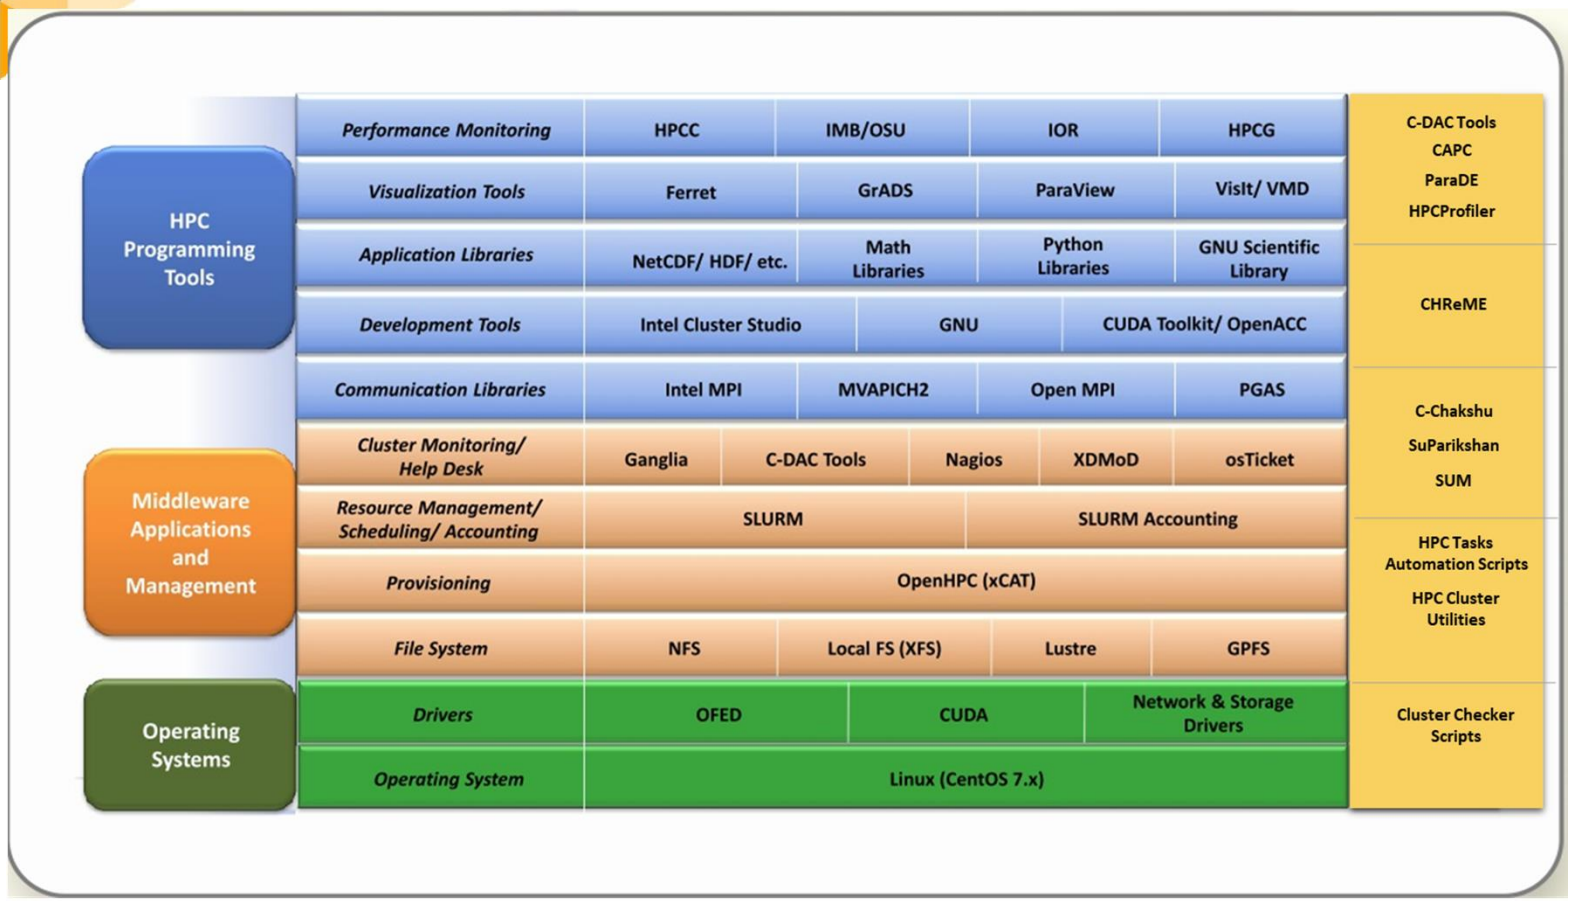

- C-CHAKSHU – Central portal
- User Creation Portal
- Accounting and billing
- Ticketing system with FAQ
- Handling more than 6000 users
- Centralized monitoring of all systems

- C-DAC's indigenous value added products

CHReME Winner of International Data Corporation Award In SC 2011, Seattle, WA

NSM Systems – Applications, Tools, Programming Models- AI & HPC

HPC Applications	
Bio-informatics	MUMmer, HMMER, MEME, PHYLIP, mpiBLAST, ClustalW
Molecular Dynamics	NAMD (CPU & GPU), LAMMPS(CPU & GPU), GROMACS
Material Modeling, Quantum Chemistry	Quantum-Espresso, Abinit, CP2K, NWChem,
CFD	OpenFOAM, FDS, SU2
Weather, Ocean, Climate	WRF, RegCM, MOM, ROMS
Disaster Management	ANUGA Hydro

Visualization Programs	GrADS, ParaView, VisIt, VMD
Dependency Libraries	NetCDF, PNETCDF, Jasper, HDF5, Tcl, Boost, FFTW
Programming Models	MPI, OpenMP, OpenACC, CUDA, PGAS, Pthreads

Installed additional applications, libraries, tools on different NSM systems as per requirements from users of respective systems

AI/ ML/ DL Tools/ Technologies	
	DL Frame work: TensorFlow , keras, theano, pytorch, scikit-learn,scipy, cuDNN
	Data Science: Numpy , RAPIDS
	Distributed DL Framework: TensorFlow with Horovod
	Container Technology: enroot
	JupyterHub: DL application development platforms and web based IDE

Application Porting, Optimization And Scaling Services

Port/tune/accelerate applications on HPC systems using recent advances in technology

Address challenges in optimization and scaling of applications on large systems

Proliferate HPC & DL through trainings, hackathons and bootcamps

Support application developers and users on NSM infrastructure

- Make available applications & tools from engineering & scientific domains and provide support
- Test, scale and benchmark applications
- Parallelize and optimize user applications
- Cater needs of end-users by adopting emerging tools & technologies including AI/ ML/ DL

HPC Applications for National Need

- **NSM Platform for Genomics and Drug Discovery**

  C-DAC, IISc, NII Delhi, IIT Delhi, NCBS Bangalore, NIBMG Kolkata

- **Urban Modelling: Science Based Decision Support Framework to Address Urban Environment Issues (Meteorology, Hydrology, Air Quality)**

  C-DAC, IISc, IITM Pune, IIT Bhubaneswar, IMD

- **Flood Early Warning and Prediction System for River Basins of India**

  C-DAC, PEC University of Technology, IISc, CWC

- **HPC Software Suite for Seismic Imaging to aid Oil and Gas Exploration**

  C-DAC, Osmania University, NGRI Hyderabad, IIT Roorkee, ONGC

NSM Platform for Genomics and Drug Discovery (NPGDD)

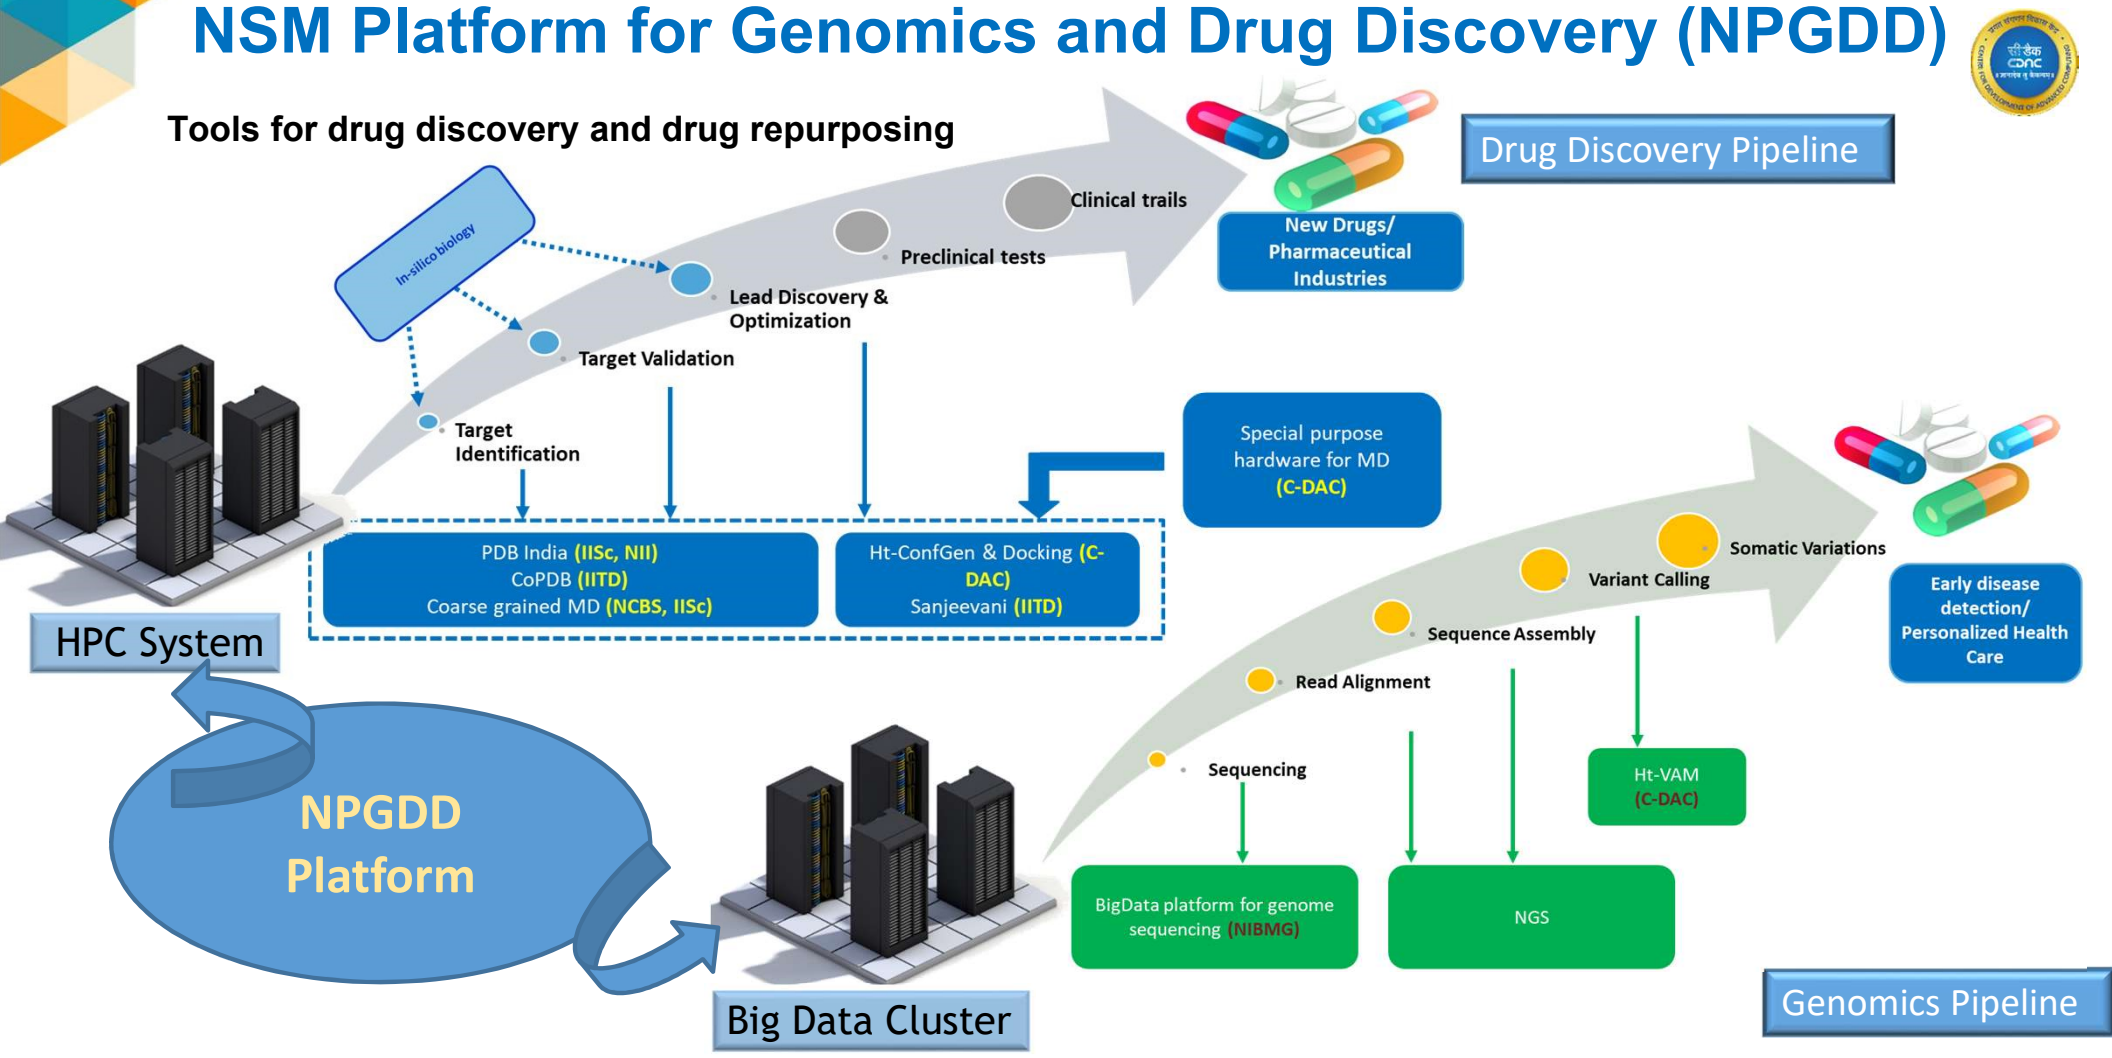

NSM Urban Modeling - Integrated System

Urban Database Management
Data localization
Interoperable data
Processed data cloud

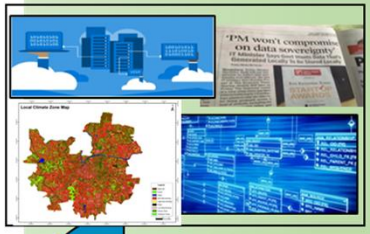

Decision Support & Dissemination System
Early warning and advisory
Data analytics and scenario generation
Weather, Flood, Pollution forecast dissemination

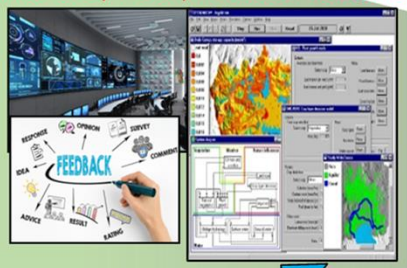

HPC enabled Automated Coupled Model Simulation Platform
Integrated cloud service for Met-Hydro-AQ-CFD Modeling and forecast generation

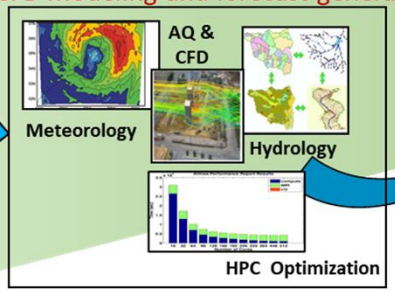

Services

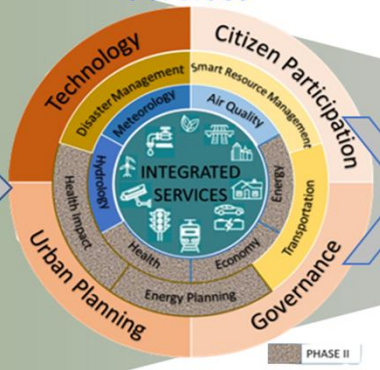

USERS

Researchers

Govt. Bodies

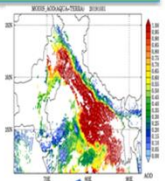

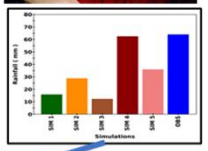

Citizens

Disaster Management Authorities

Central & State Pollution control boards, Municipal corporations

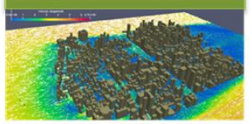

Partners

C-DAC IMD IISc IITM Pune IIT Bhubaneswar

Early Warning System For Flood Prediction

- Pilot is for Mahanadi River Basin
- 3-days flood forecast
- Inundation (submerged) and Water level forecast
- **End Users:** Central Water Commission (CWC), State Water Resources Department, State Disaster Management Authority, District Administration
- Forecast is being shared with CWC for last 2 years for verification

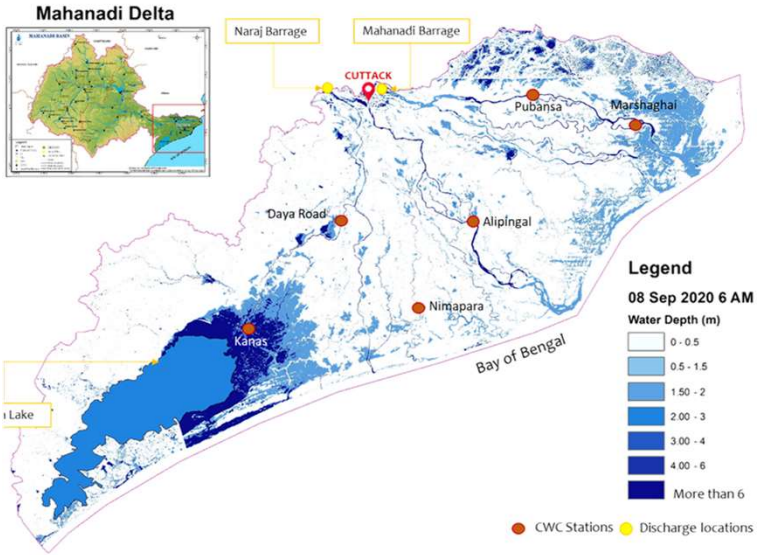

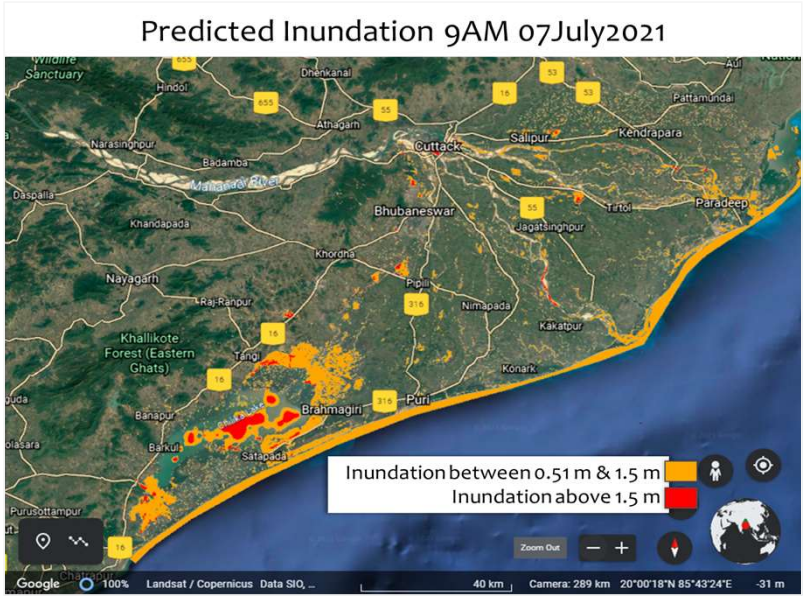

Simulation outputs. a: Mahanadi Delta simulation output; b: Simulation output overlaid on base map

NSM Applications: Seismic Imaging Software

HPC Software Suite for Seismic Imaging to Aid Oil & Gas Exploration (SeisRTM)

End Users

- Agencies involved in oil and gas exploration - ONGC
- Research organization for deep crustal studies
- Academia for teaching advance seismic processing

Collaborators

- ONGC
- IIT Roorkee
- CSIR- NGRI
- Osmania University

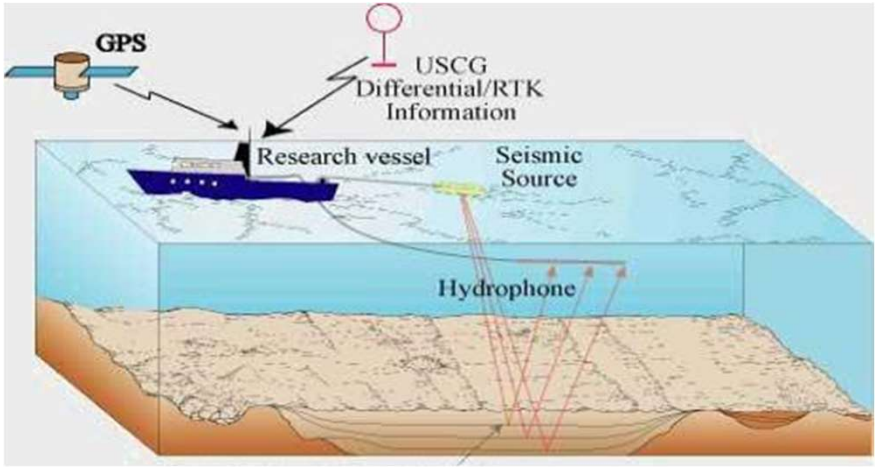

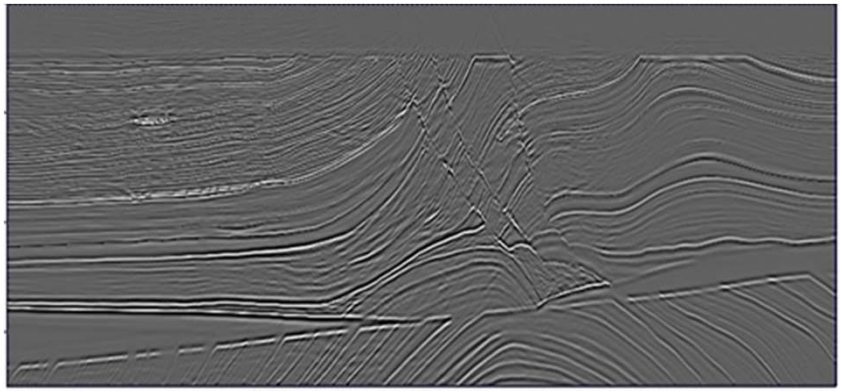

Subsurface Earth Structure up to 12 kms below ocean floor

Creation of 20000 HPC aware manpower under NSM

- As on date, **19500** HPC manpower trained by involving academia and industry
- Types of Courses
  - Short term – Workshops, bootcamps, hackathons
  - Medium term – PG Diploma courses, internships
  - Formal courses – conducted in educational institutes
- NSM Nodal centers at
  - IIT Kharagpur, IIT Madras, IIT Goa and IIT Palakkad
- Collaborating Industry partners
  - Intel, Nvidia, OpenACC, Ansys

Some of the Initiatives on System Software, Management, benchmarking and Application software

C-CHAKSHU: Multi-Cluster Monitoring Platform

- **Live Dashboard**: Live data in a single page to view several aspects about the overall operation of Multiple HPC system
- **Monitor:** Node, Infrastructure, Jobs, Users, CPU, Services, Interconnects
- **Analyse:** Plan More Intelligently by forecasting usage based on real, historical usage data
- **Time-frame based Job Accounting and Reporting:** Plan more intelligently by forecasting usage based on real historical data
- **Health Monitoring:** Drill down Nodes, Interconnects view and alerts

C-CHAKSHU: Multi-Cluster Monitoring Platform

- **HPC Usage Billing:** Chargeable HPC recourses usage billing as per project, user type
- **3D Rack View:** Quick identification of issues in nodes, network, storage
- **Integrated with Ticketing System**: Better insights to user support request and subsequent actions by administrator, decision maker
- Support for industry standard schedulers like SLURM, LSF
- Poster Selected and Presented in ISC 2023
- Real-time Application Performance Monitoring
- CPU and Memory representation in Time series data

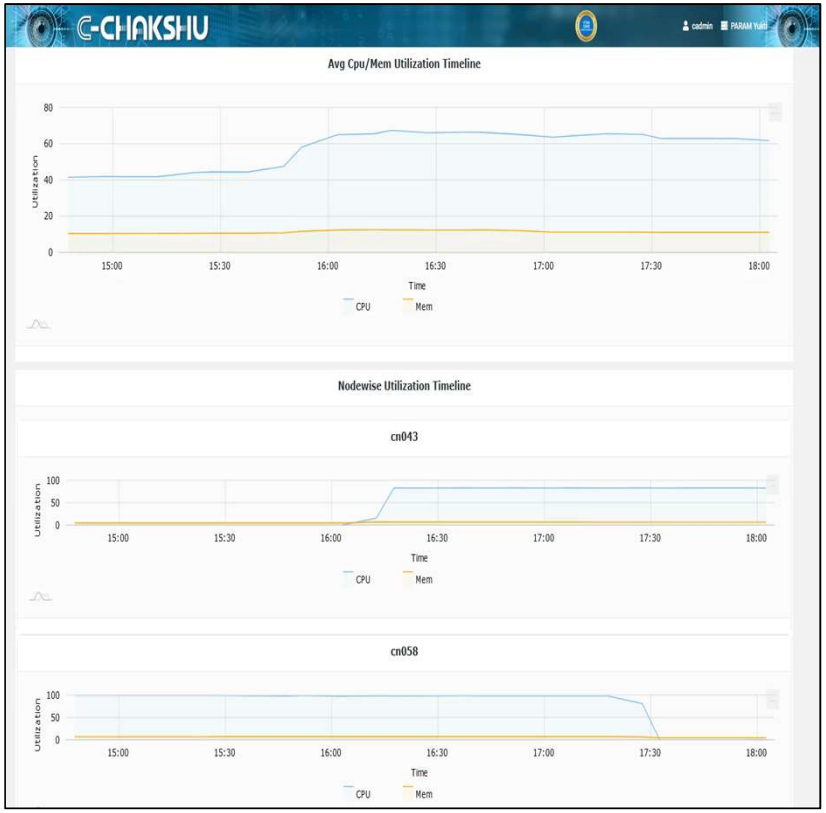

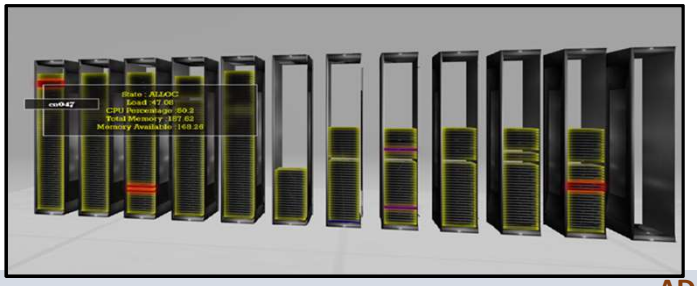

C-UPE

(C-DAC Unified Programming Ecosystem)

Developing the complete software ecosystem to support Unified programming using SYCL for C-DAC AUM HPC Processor

- C-SYCL : Developing SYCL compiler for CDAC AUM HPC processor
  - ✓ Creation of SYCL header files created
  - ✓ Developing clang based compiler frontend for SYCL
  - ✓ Developing SYCL support for compiler backend
- SYCLProf : Support profiling of application using C-SYCL
  - ✓ SyCL application profiling using TAU – with OpenSyCL and CUPTI libraries
- SYCLDebug
  - ✓ Development of SyCL application debugging modules for GDB and LLDB on ARM architecture
- SYCLLib
  - ✓ Development of SYCL aware libraries which support C-SYCL compiler on ARM architecture

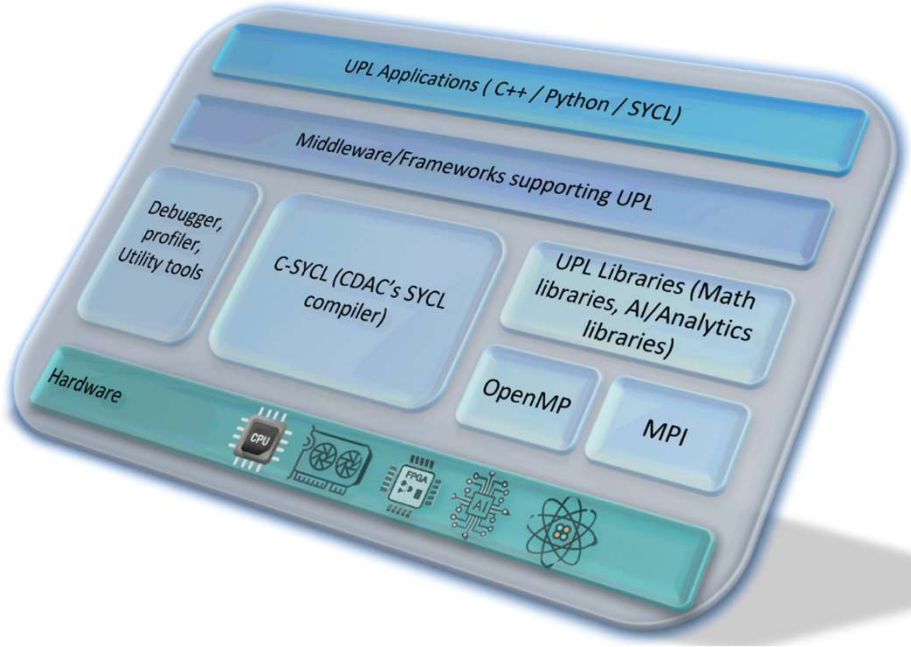

Compiler for C-DAC HPC Processor

Objectives:

- ✓ Optimization of GCC compiler for C-DAC AUM HPC processor
- ✓ Development of C-DAC compiler based on LLVM for C-DAC AUM HPC processor

Activities being carried out

❑ GCC Optimization

- **Design and development of a Performance library-based Plugin**
  - ✓ Use of optimized math libraries to speed up the execution of linear algebra operations
  - ✓ Observed 8X-9X performance boost for GEMM
- **Development of Plugin to control Loop unrolling factor**
- Optimizations for increasing the vector register count in GCC, restrict the number of branch instructions, cache missed and branch mis-predictions
- **Development of plugin for Memory optimization**

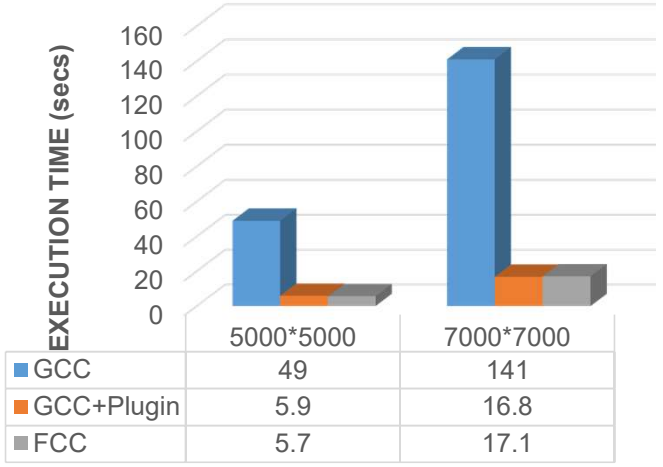

Performance Evaluation with GEMM(Double)

Compiler for C-DAC HPC Processor (contd..)

❑ Development of C-DAC Compiler based on LLVM

- Integration of external vectorizer with LLVM
  - ✓ Observed 5-7 times better performance for GEMM
  - ✓ Significant performance improvement in Tealeaf and miniBUDE
- Created and deployed a code patch which adds functionalities for modification in unroll factor and interleave count
  - ✓ Observed 2 times performance boost for GEMM
- Modification of LLVM Cost model to support outer loop vectorization
- Enabling SVE support for outer loops
- Optimizations for reducing the number of stall cycles, cache misses, branch mispredictions etc.

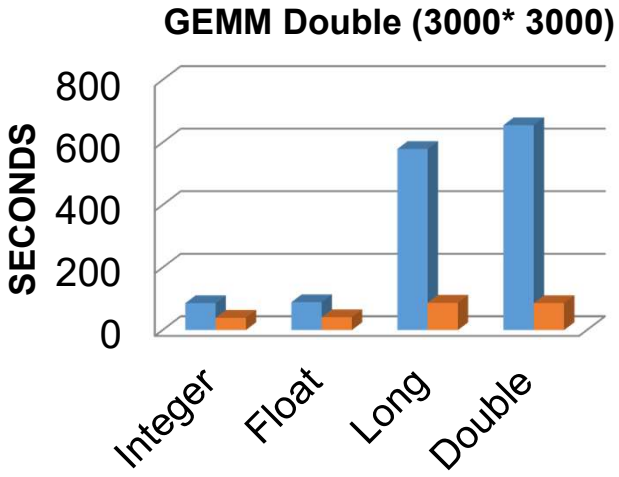

Performance with External Vectorizer

Libraries (MPI – Math – AI/ML)

- MPI: Optimization of OpenMPI for CDAC AUM HPC processor

Activities being carried out

- Holistic comparison of different platforms
- Enhancing OpenMPI performance on ARM
  - ✓ Enabling vectorization(SVE)
  - ✓ Code optimization

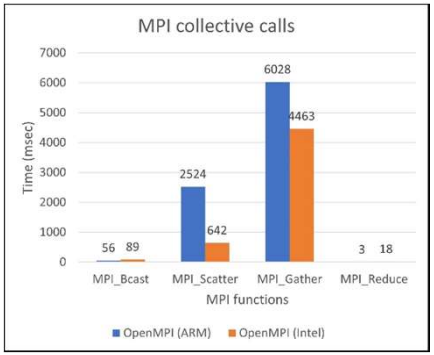

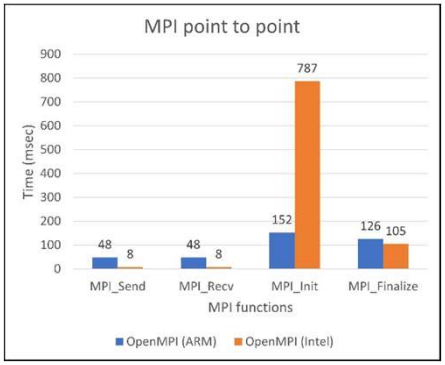


```
(venv) (base) [deepika@master01 my-python-project]$ pip3 freeze
absl-py==1.2.0
arm-mango==1.2.0
astunparse==1.6.3
attrdict==2.0.1
cachetools==5.2.0
caffe-ssd-x86==1.0.0.1
certifi==2022.6.15
chainer==7.8.1
charset-normalizer==2.1.1
click==8.1.3
cycler==0.11.0
dataclasses==0.6
etils==0.7.1
exchange-calendars==4.2
filelock==3.8.0
flatbuffers==1.12
fonttools==4.36.0
frozendict==2.3.4
gast==0.4.0
google-auth==2.11.0
google-auth-oauthlib==0.4.6
google-pasta==0.2.0
graphviz==0.8.4
grpcio==1.47.0
h5py==3.7.0
idna==3.3
imageio==2.21.1
importlib-metadata==4.12.0
importlib-resources==5.9.0
jax==0.3.16
joblib==1.1.0
keras==2.9.0
Keras-Preprocessing==1.1.2
kiwisolver==1.4.4
```

- **AI/ML Libraries/packages :** porting and optimization of HPC based AI/ML libraries on ARM architecture
  - ✓ Porting of the popular libraries & packages on ARM
  - ✓ Created a single image for installation
  - ✓ Benchmarking & optimization of the identified

Libraries (MPI – Math – AI/ML) contd..

Math Libraries: Optimization for BLAS , LAPACK FFT for CDAC AUM HPC Processor

Activities being carried out

- Benchmarked Math Libraries
- Enabling vectorization using ACLE for BLAS

One-Dimensional (DP) without Optimization:

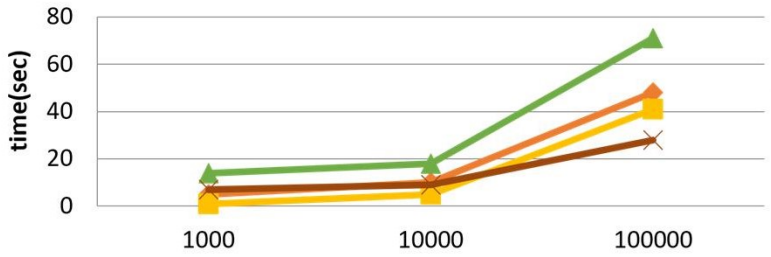

One-Dimensional (DP) with Optimization:

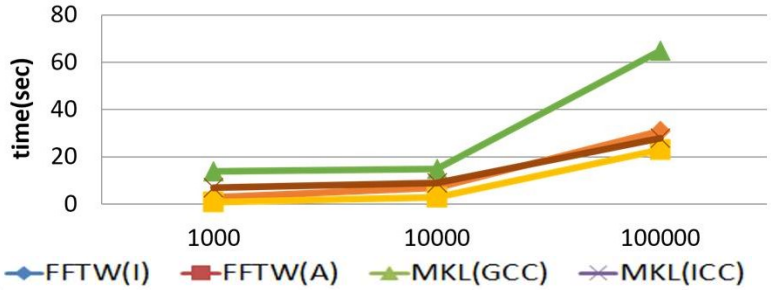

Matrix-Matrix Multiplication without Optimization:

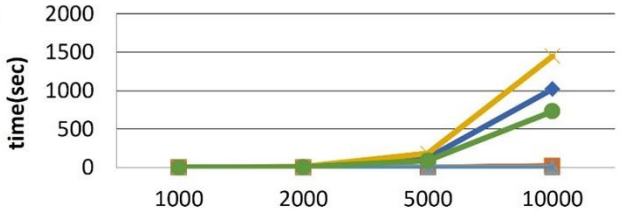

Matrix-Matrix Multiplication with Optimization:

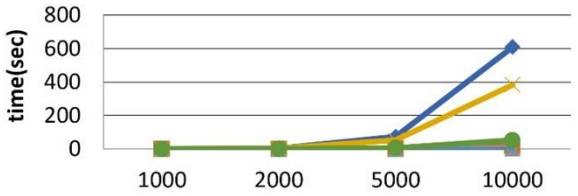

Linear Equation without Optimization:

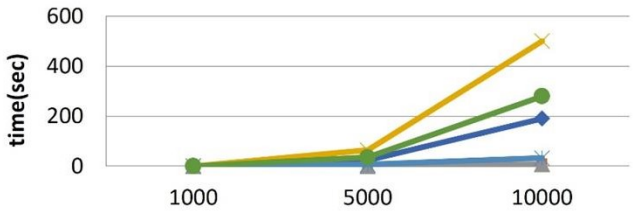

Linear Equation with Optimization:

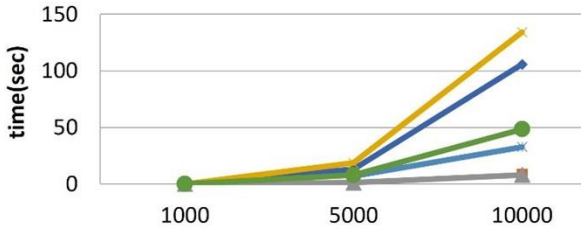

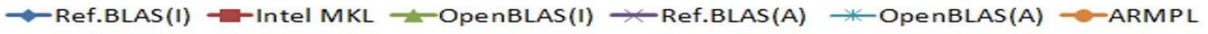

Real-time processing of radio telescope, GMRT data with HPC and AI

NCRA (TIFR) and C-DAC

- GMRT will be the only facility equipped to provide simultaneous detection and arc-sec localization over the full 300−1460 MHz radio spectrum preparing for the Square Kilometer Array (SKA)
- Discovering 250+ of Fast Radio Bursts with the GMRT in 3 years science operation by piggy-backing on existing observations (~ 10-times increase of current known population created after a decade of operation!)

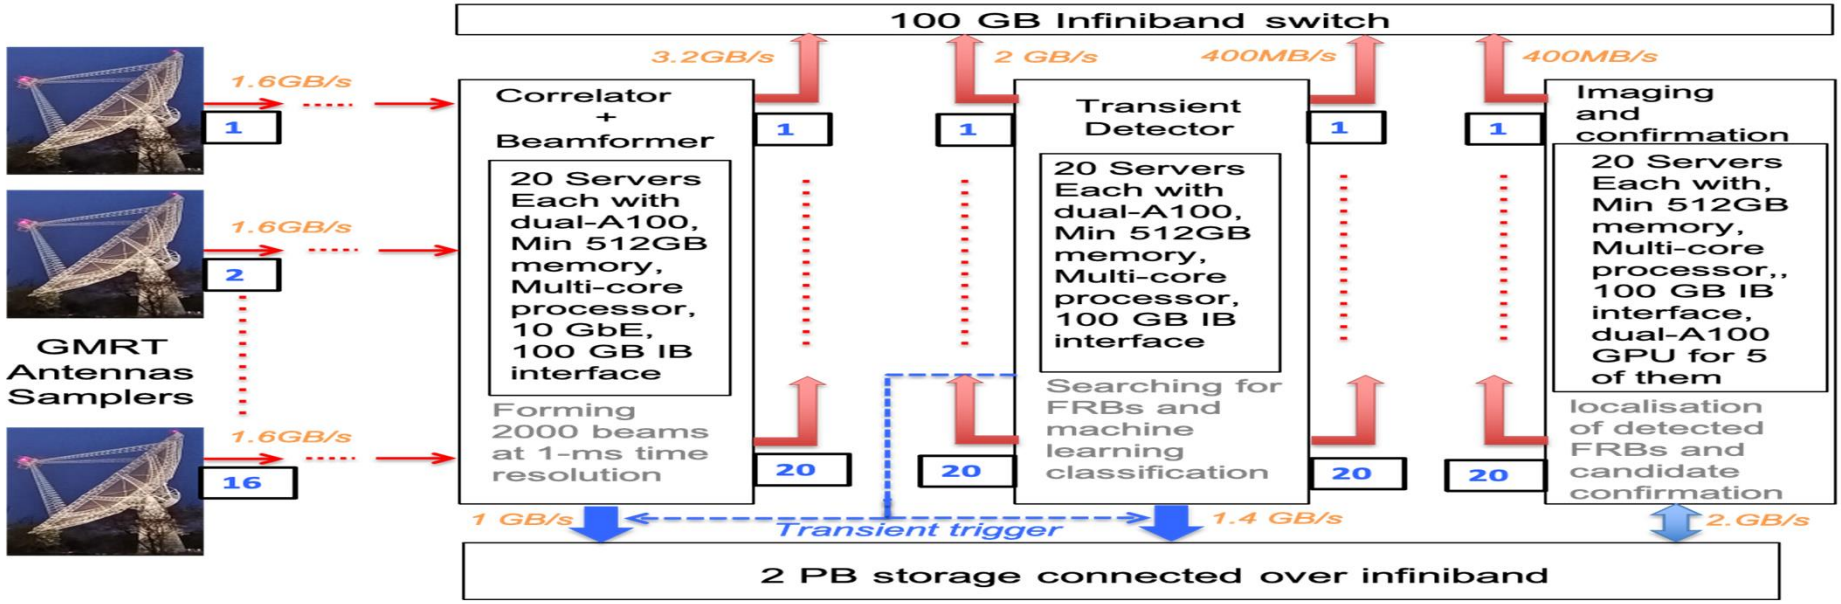

Application Optimization

ANUGA: Optimization of delta simulation through finer mesh resolution

- ANUGA is a free and open source software tool for hydrodynamic modelling, suitable for predicting the consequences of hydrological disasters such as riverine flooding, storm surges and tsunamis.
- Optimized the code and achieved a performance gain of approximately 6X with 900 sq.m resolution
- The code was also toned down from 8100 sq.m resolution to 300 sq.m

2 Days simulation : 300 sq m res. ; 15th Sep ; MAH=5mm MAS=0.1 CFL=1.7		
Nodes	MPI Ranks	Time
64 (46C/N)	2944	6.0 Hrs
96 (46C/N)	4416	4.6 Hrs

Mesh Resolution (sq. m)	Maximum RSS Slave Process	Maximum RSS Master
900	0.82 GB	42.03 GB
300	2.86 GB	118.8 GB
100	6.33 GB	325.2 GB

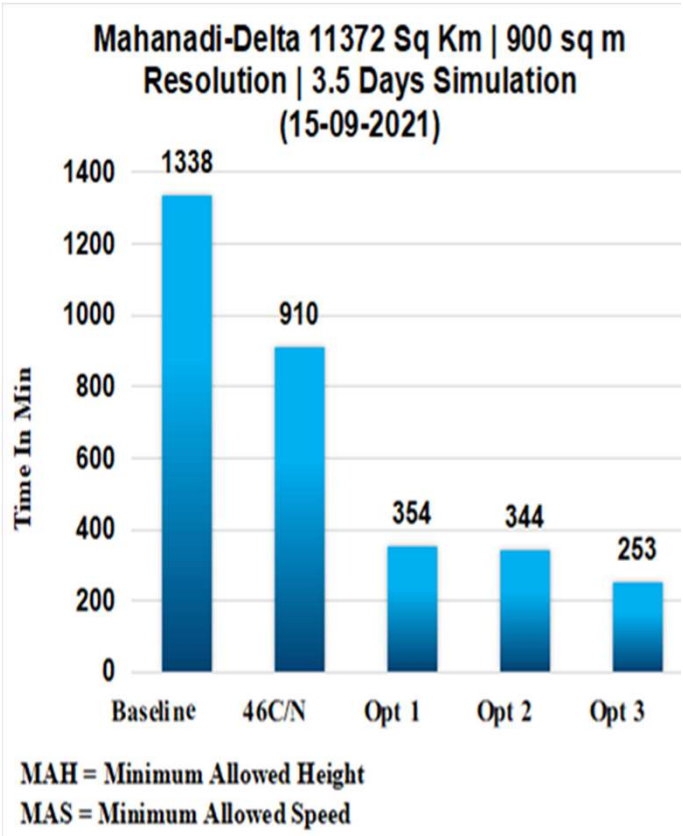

Application Optimization

ANUGA: Work in Progress

Memory Optimizations

- Memory access pattern for different resolutions
- Optimizing ANUGA for higher mesh resolution
- Choosing single precision over double precision wherever possible

GPU porting:

- Using CUDA C and CuPY kernel for compute intensive code
- Using pragma based GPU offloading like OpenACC and OpenMP

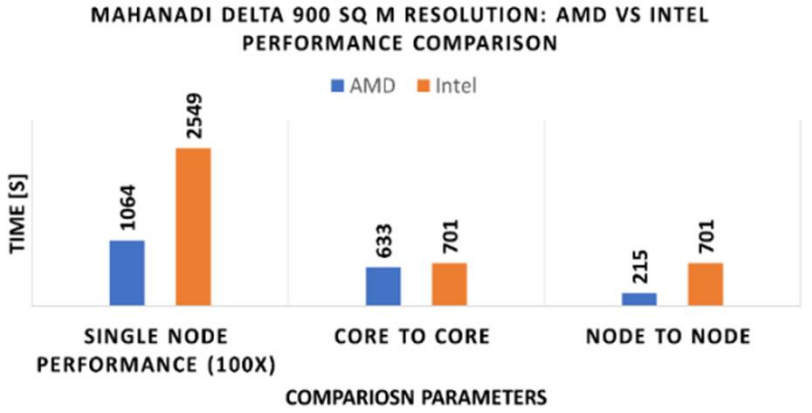

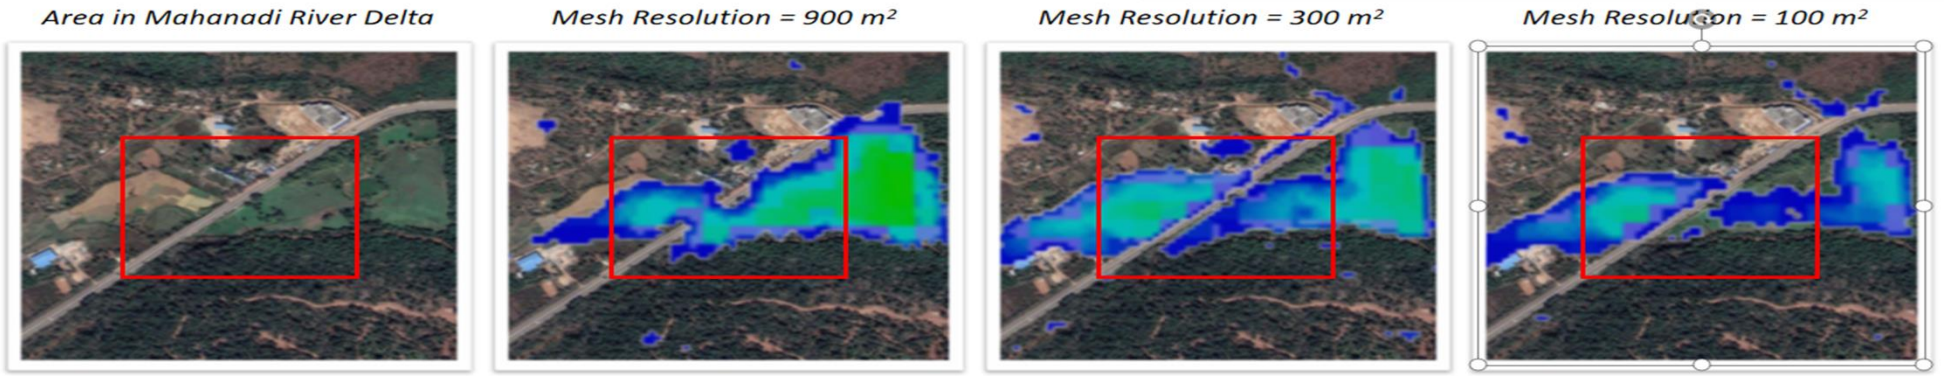

Application Optimization

TRITON: Flood simulation for Exascale Machines

C-DAC & IIT Bombay

- **Two-dimensional Runoff Inundation Toolkit for Operational Needs**
  - It can operate on multiple computer platforms and utilize modern HPC environments.
  - Implementation with a single central processing unit (CPU) or multiple CPUs (using OpenMP+MPI)
  - Implementation with a single graphics processing unit (GPU) or multiple GPUs (using CUDA+MPI)
  - Highest TRITON computational efficiency can be achieved by using GPU implementation
  - TRITON utilizes topographical data (e.g., digital elevation model [DEM]

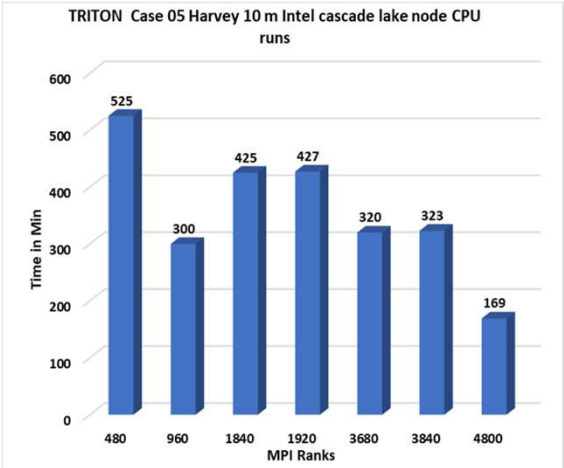

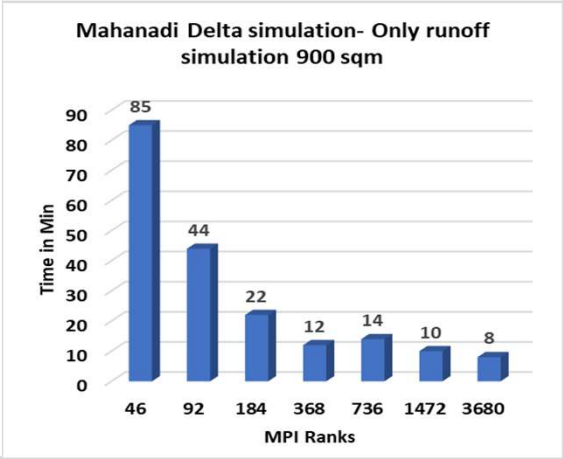

Exploration of Upcoming Technologies

Oil Exploration: SYCL

- **Objective:** To port CUDA based SeisAcomod-2D to Intel oneAPI DPC++ (SYCL) to make it runnable on different architectures
- **Accomplishment:** Ported CUDA based open source SeisAcoMod2D code to DPCPP (SYCL) followed by validation of results from domain experts. This enabled single unified source code to be run on Intel Data center GPU Max and Nvidia A100 GPUs along with Intel CPU.
- **Future Explorations:** To test the migrated SYCL code on other compute platforms like AMD, FPGA etc.

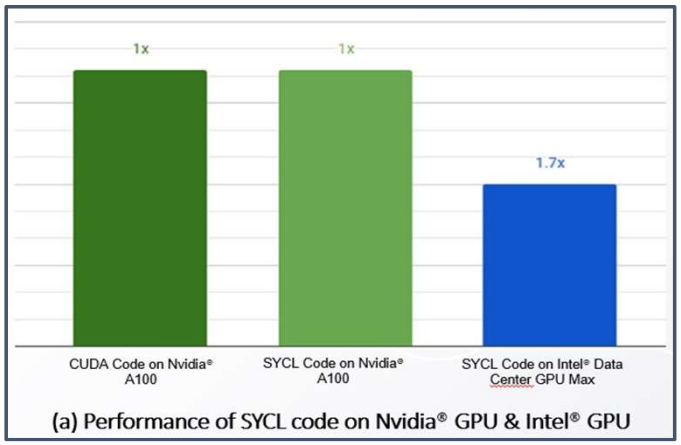

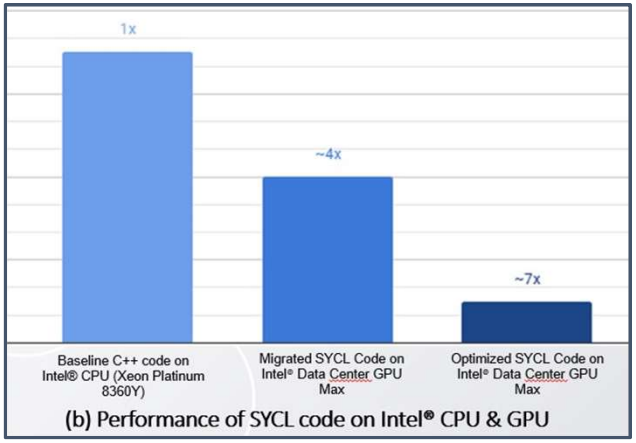

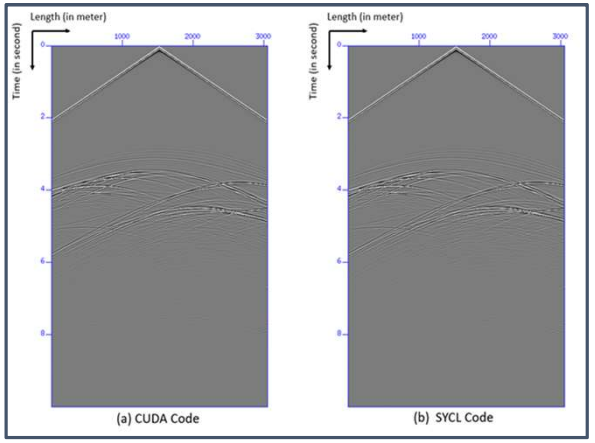

Exploration of Upcoming Technologies

KOKKOS

- **Objective**: To explore the LAMMPS application performance with and without KOKKOS feature
- **Platforms**: ARM A64FX (CPU), Intel Cascade Lake (CPU), NVIDIA-A100 (GPU), AMD Instinct MI210 (GPU)
- **Accomplishments**: We have tested an open source application LAMMPS with & without KOKKOS on different compute platforms and observed it's performance.
- **Future Plans**: Enablement of other HPC applications with KOKKOS on these platforms

Application: LAMMPS							
ARM A64FX (CPU)		Intel Cascade Lake (CPU)		NVIDIA-A100 (GPU)		AMD-MI210 (GPU)	
Baseline	With KOKKOS	Baseline	With KOKKOS	Baseline	With KOKKOS	Baseline	With KOKKOS
1217 m	810 m	311 m	226 m	4m 7s	3m 42s	6m 25s	6m 26s

Explorations of Different Compute Architectures

AMD GPU vs Nvidia GPU Benchmarks on RUDRA

Objective: Examine AMD Instinct MI200 Series accelerators for HPC and DL workload.
Platform Configurations: Rudra, CPUs: 48, CPU Model: Intel(R) Xeon(R) Gold 6240R CPU @ 2.40GHz, GPUs Used: NVIDIA A100 (PCIe 80 GB) & AMD Instinct MI2100 (PCIe 64 GB)

Observations:

- **ML/DL:** (TF-CNN) Nvidia A100 shows better perf. as compared to AMD-Instinct MI210 with **1.34x** performance on **2** GPU cards
- **NAMD, CP2K, SpecFEM-3D :** NAMD & CP2K application shows **comparable** performance on both NVIDIA A100 & AMD Instinct. SpecFEM-3D performs 1.78x better on Nvidia platform

Future Plans: Performance analysis of the AMD GPU platform with other vendors GPU platforms wrt to different codes / applications

Fujitsu's ARM A64Fx Explorations

Objective: Explore the Fujitsu's ARM A64FX compute architecture for different scientific domain HPC applications.
Platform Configurations: Fujitsu A64FX (48cores, 1.8GHz)

Applications: NAMD, LAMMPS, GROMACS, OpenFOAM, WRF, SpeckFEM-3D, Quantum Espresso, Abinit, CP2K, Tensorflow etc.

Observations:

- Explored HBM2 capability of ARM clusters through use codes and applications with large data size
- Benchmarked memory intensive codes/ applications on ARM & other HPC clusters.

Future Plans: Microarchitecture level explorations, Performance analysis & benchmarking of other codes/ applications

Exploration of Upcoming Technologies: Julia

Objective :

- Performance evaluation with C
- Scalability and stability examination
- Scientific applications suitability

Test Bed Configurations:

- CPU Model: Intel(R) Xeon(R) Gold 6148 CPU @ 2.40GHz, CPUs: 48
- GCC -12.2.0, OpenMPI-4.1.2, Python-3.8, Julia-0.19.0

Observations:

- Better performance than Python
- Comparable performance than C
- Development is faster & easy as in the Python

Future Activity

- Peta and Exa Scale applications exploration
- ML/DL application exploration

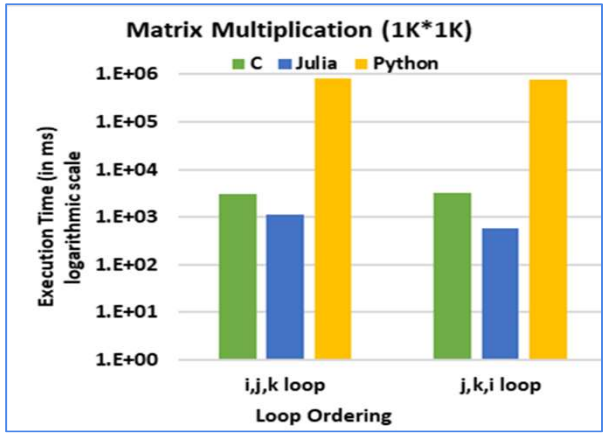

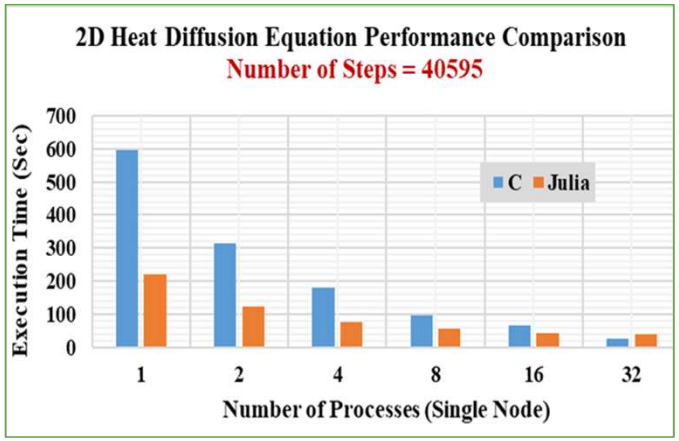

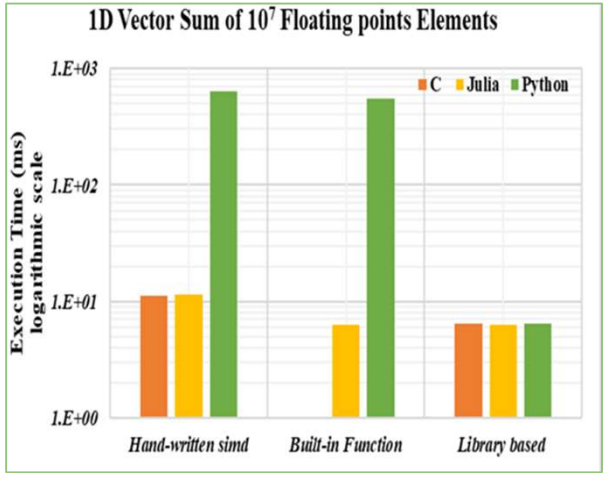

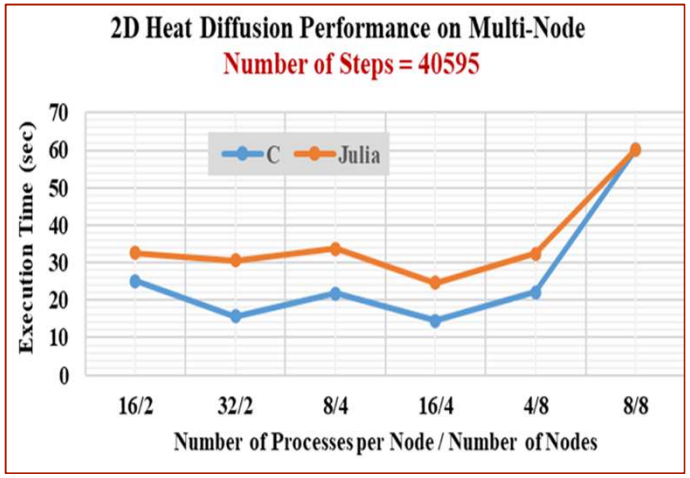

Exploration of Upcoming Technologies

Finance Domain Exploration

HPC Uses in Finance Domain

- Credit risk assessment, Cyber security, Fraud detection, Algorithmic trading, Strategy development process in algorithmic computing

Future Activity

- Scaling existing prediction models across HPC Cluster
- Porting C/C++ models into Python using pyomp library for scalability across multiple CPU cores

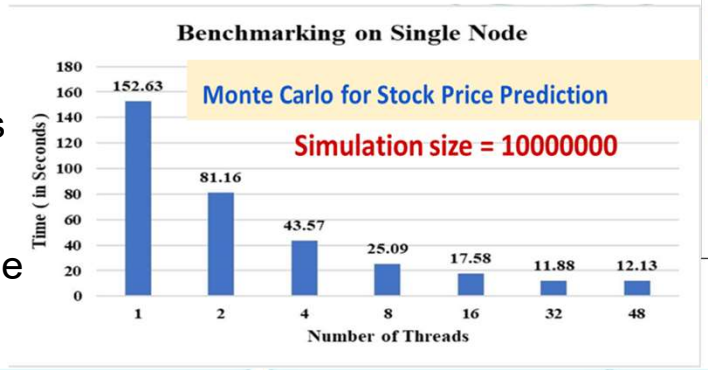

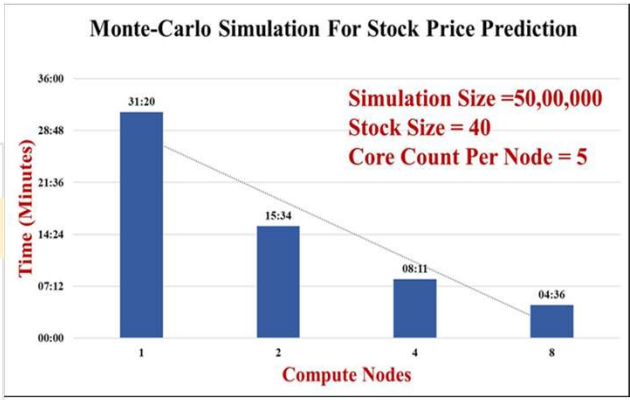

Finance Domain Software Stack

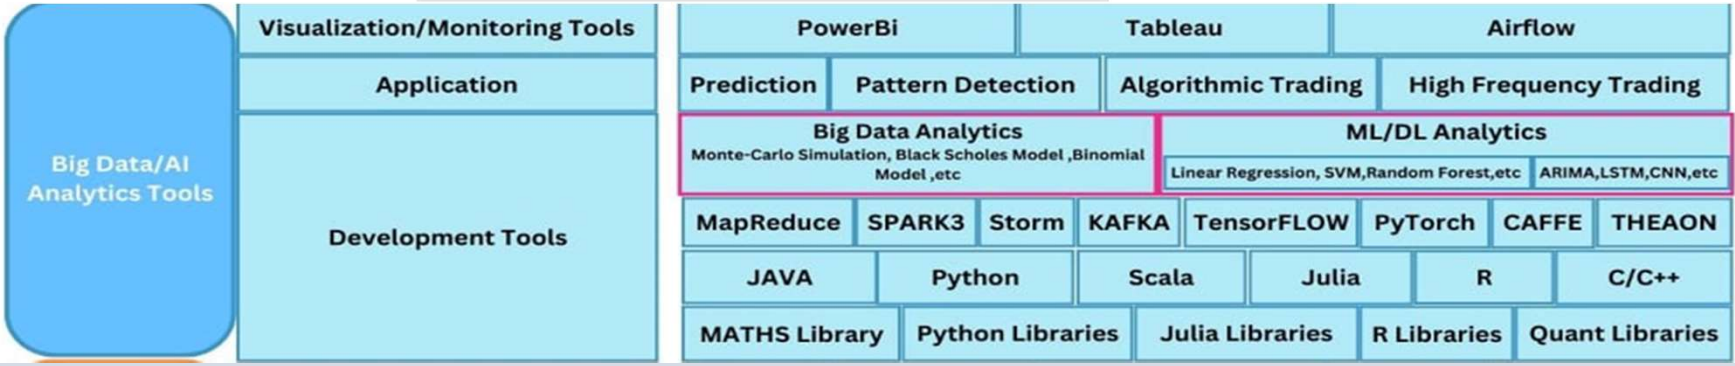

AI/ML/DL on NSM

Expertise in HPC-AI applications across heterogeneous & hybrid platforms
Porting, Optimizations & Scaling of AI applications (Mutli-Node Multi-Accelerator setup)
Interested in working on Traditional Scientific Simulation Applications and new age HPC-AI Converged Application life cycle

Tools/ Technologies/ Libraries	**DL Framework:** TensorFlow , keras, theano, pytorch, scikit-learn,scipy, cuDNN	
	Data Science: Numpy , RAPIDS	
	Dist. DL Framework: TensorFlow with Horovod	
	Container Technology: enroot	
	DL App. Dev. Platform, web based IDE: JupyterHub	
R&D	Automatic Number Plate Recognition	ANPR for campus security
	User Behavior Analysis	**Identification of malicious user login activities on NSM clusters**
	Code Identification	**Identification of malicious code execution patterns on HPC clusters**
Products	PARAM Shavak DL GPU	
	PARAM Shavak VR	

Benchmarking & Profiling	Resnet 152 Inference	RUDRA server cloud VM AWS cloud VM.
	Resnet50 Training Benchmark	RUDRA Server PARAM Siddhi ARM DevKit
	Resnet50 Training and Inference Benchmark	Graphcore IPU
	Resnet 50	Nvidia V100 GPU
MLPerf	Resnet 50 Training	Multiple node multi GPU run
	MaskRCNN Training	
	SSD Training	
	DLIO	
	Cosmoflow	
Training & Workshops	Rapids / Jupyterhub	
	Deep Learning	
	Python (pandas, numpy, matplotlib, etc.)	
	Machine Learning	
	Data Preprocessing	

Overview of Bioinformatics Activities

Supercomputing Clusters

To provide high end dedicated computing facility to researchers

Bioinformatics Resources and Application Facility (BRAF)

Computational Genomics

Mycobacterium Genomics

Population Genomics

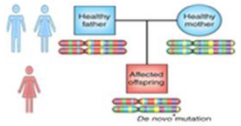

Software Development (BigData/AI/Cloud)

Molecular Modelling

Cancer proteins

Protein misfolding

Antisense Technology

Membrane Proteins

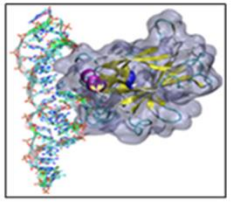

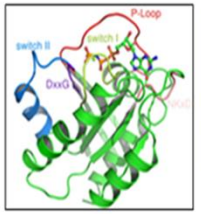

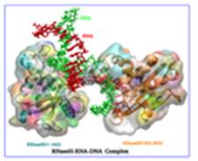

AnvayaNGS: A workflow for NGS Analysis

AnavyaNGS: Key Features

- 7 Pre-defined Workflows for frequently used pipelines in next generation sequencing.
- Rules Engine which adds intelligence to control tools connectivity.
- Token(JWT) Based Security system integrated with Spring Security.
- Intermediate tools output can be downloadable
- FTP functionality support.
- Easy to use, standalone Anvaya Client which is supported on Windows as well as Linux.

AnavyaNGS : Predefined Workflows

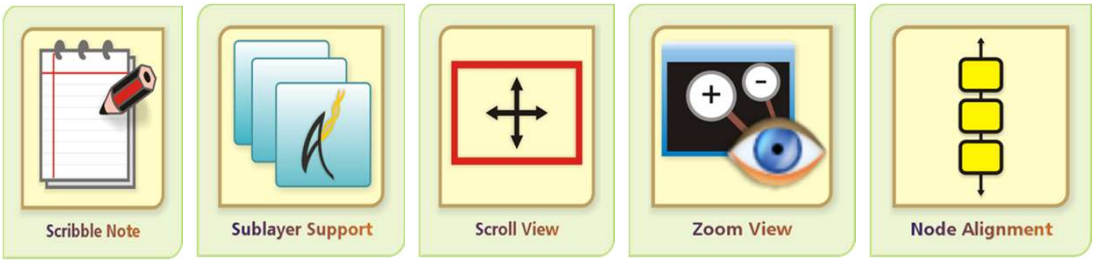

High-throughput conformation generation and docking : TANGO-Dock

TANGO-Dock Features:

- MPI based conformational search of small molecules
- Automation of protein pre-processing for Docking
- Simultaneous docking of multiple conformations of multiple ligand molecules

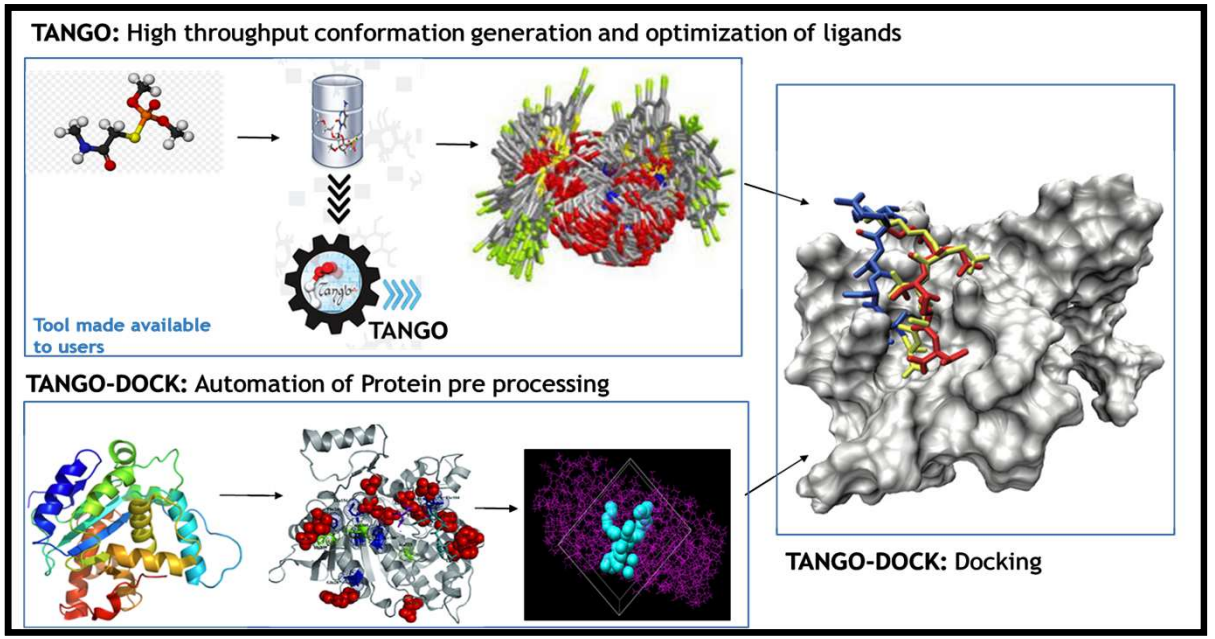

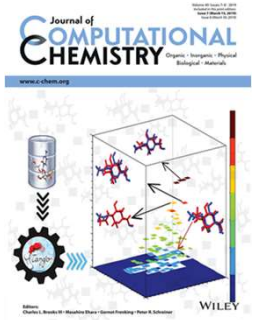

Codes like NWChem may add value to TANGO based conformational search

DPICT: Advanced Molecular Dynamics Visualization and Analysis Tool

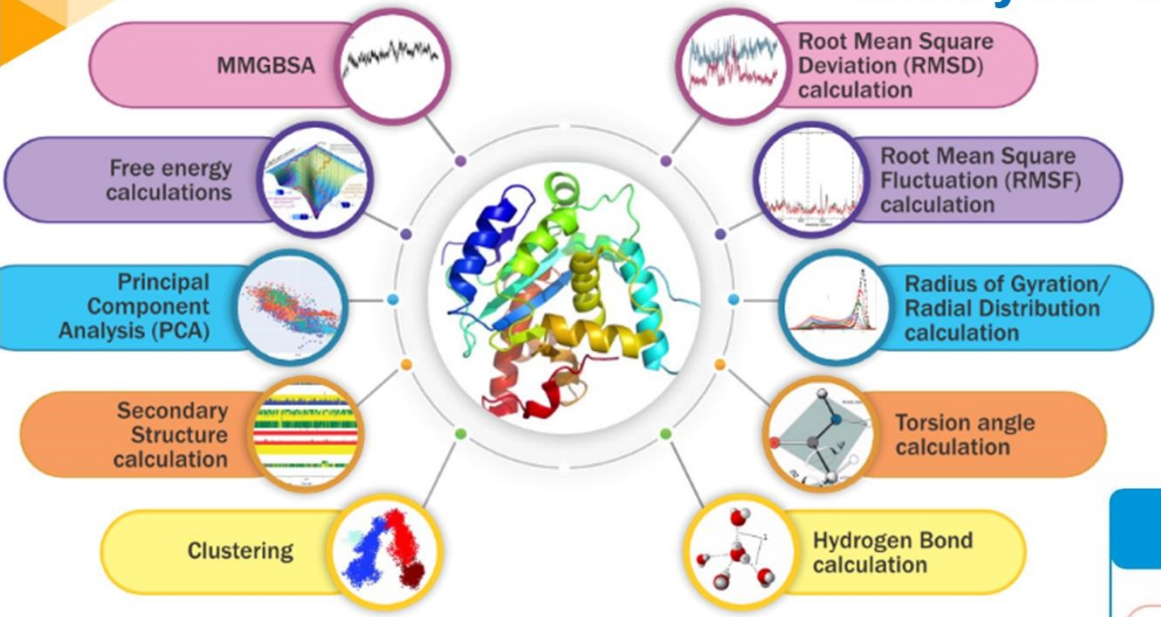

Analysis tools available in DPICT

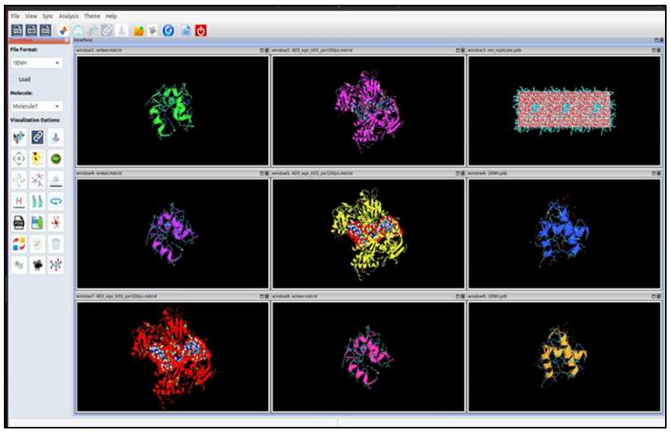

Parallel visualisation of multiple trajectories

DPICT
A parallel visualization and analytics engine for biomolecular trajectories

Deployed at

Regional Centre for Biotechnology, Faridabad

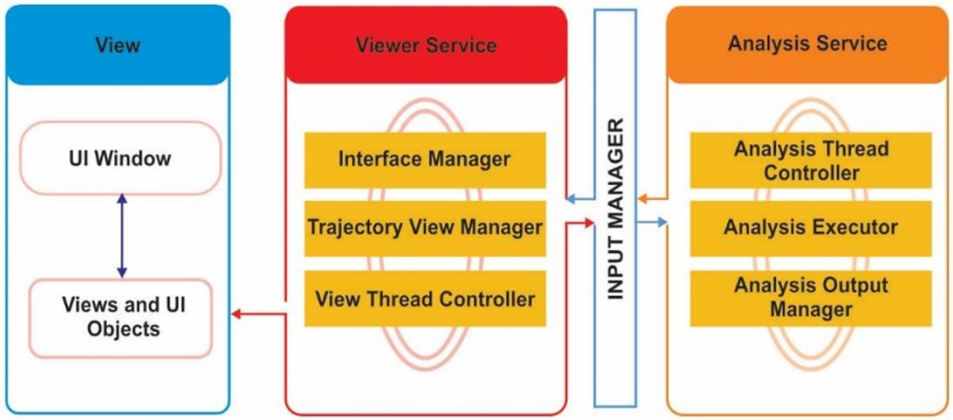

Architecture of DPICT

Development of CIMULATE: A light weight Molecular Dynamics simulator

- Indigenously develop Molecular Dynamics (MD) code
- To simulate large biomolecules in relevant time scale
- Design to tune on modern hardware
- Efficiently Scalable on CDAC's HPC environment
- Try to speed up MD calculations than existing codes

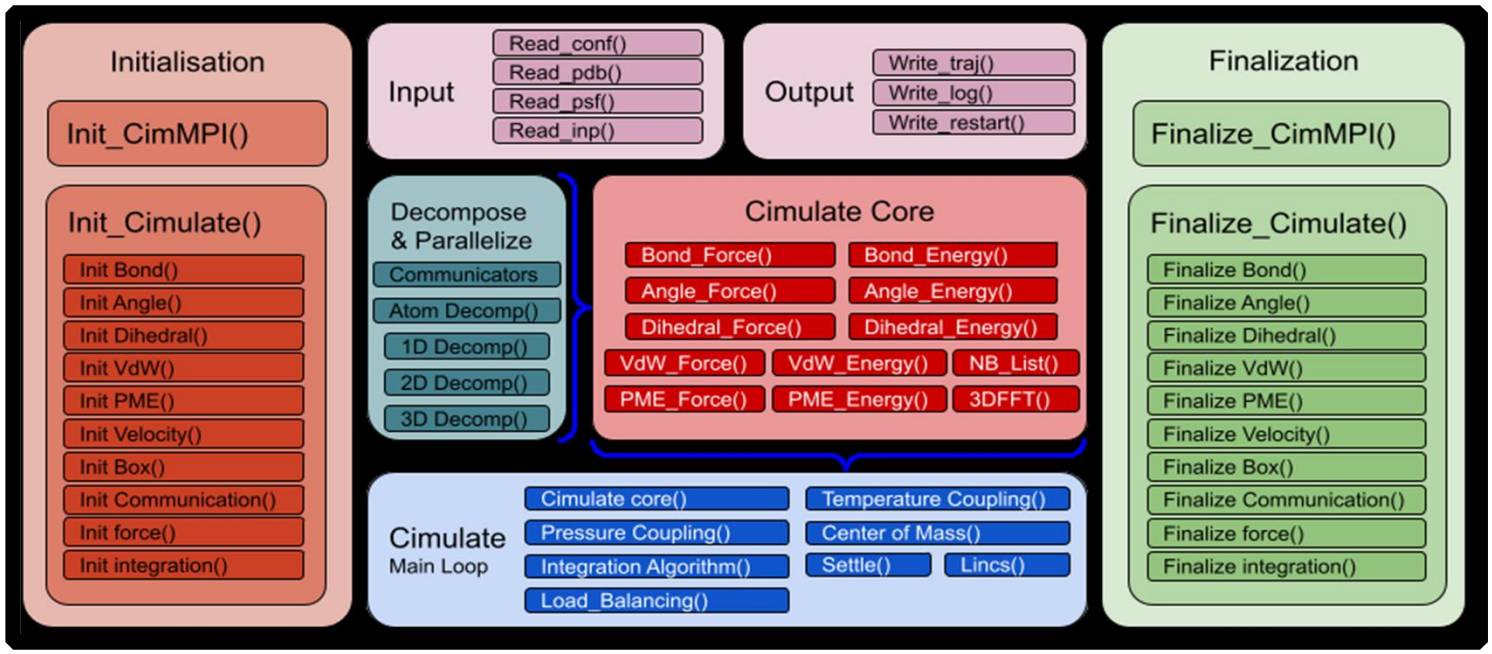

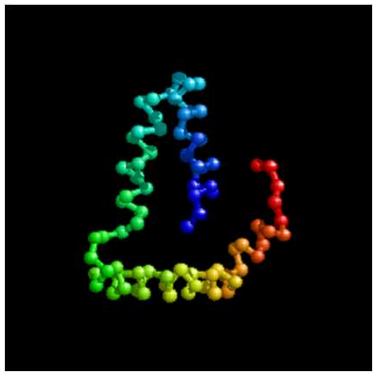

Special Purpose Computer (SPC) for Molecular Dynamics (Design & Development of a SPC for MD simulation)

- **Architecture Exploration Phase:** To arrive at the over-all System Architecture of the Special Purpose Computer for Molecular Dynamics Simulation

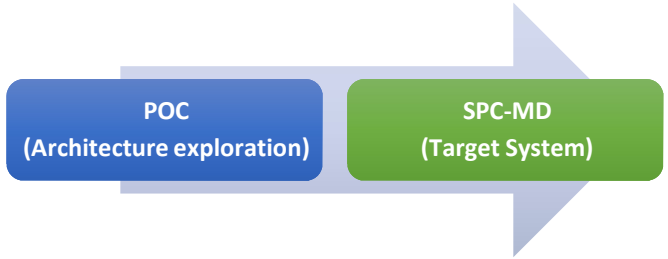

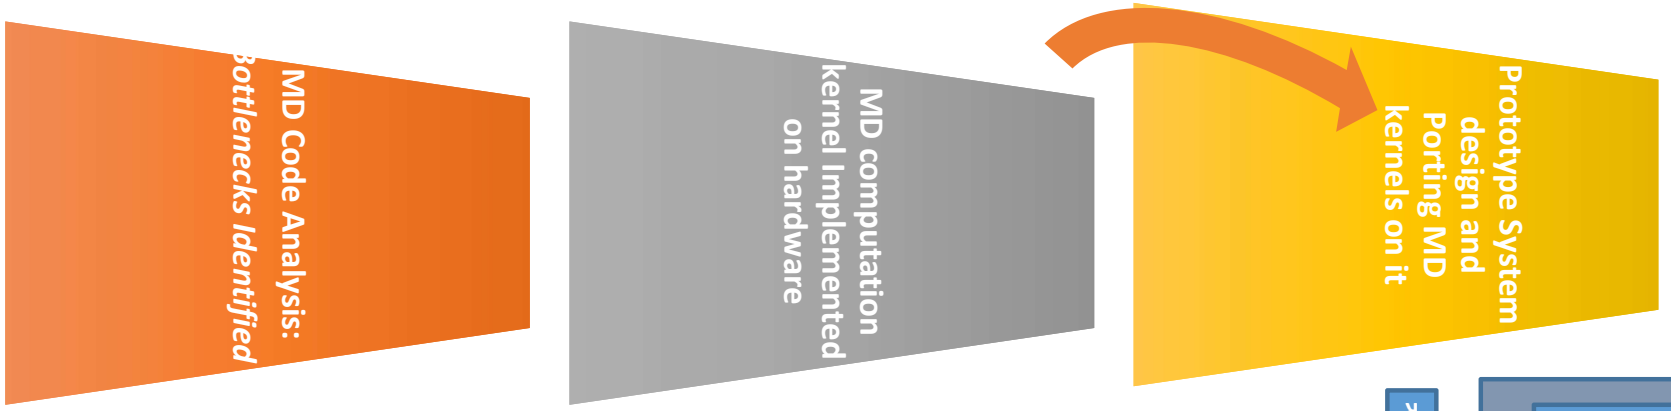

·Gromacs and Genesis code analyzed.
·Non bonded force calculation identified as bottleneck

·Non bonded (L J) computational block and 3D FFT implemented on FPGA

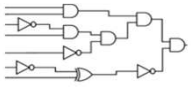

·Prototype system design:

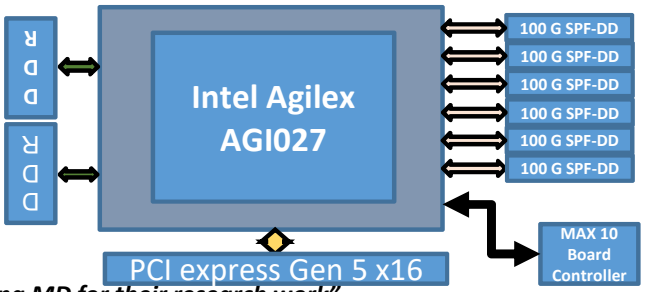

"Beneficiaries of this interdisciplinary indigenous hardware work will be the leading scientists using MD for their research work"

HPC based Drug Repurposing: COVID19 Studies

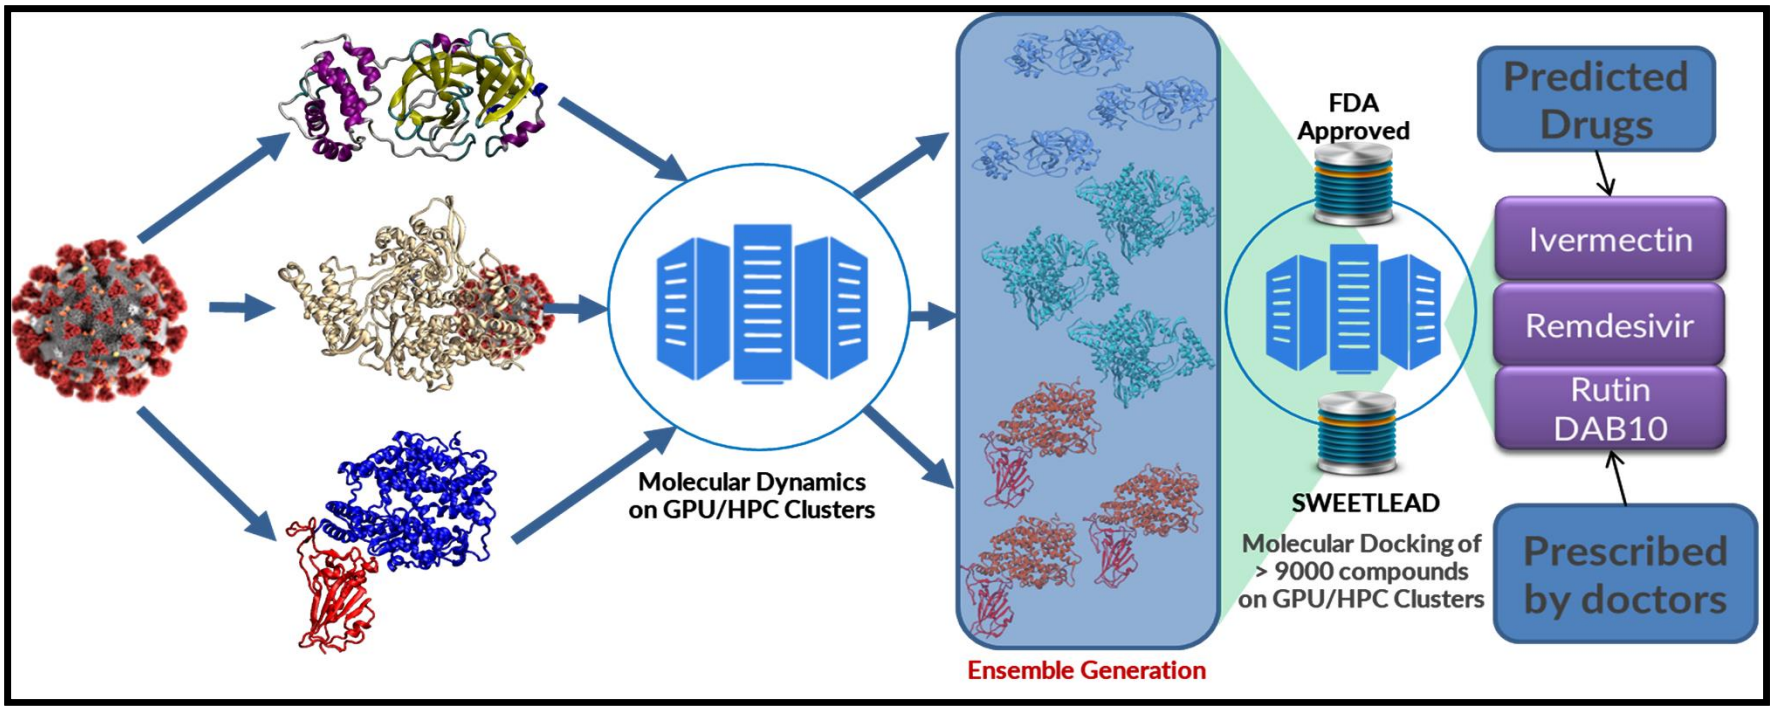

HPC based Drug Repurposing for Ayurvedic molecules: COVID19 Studies

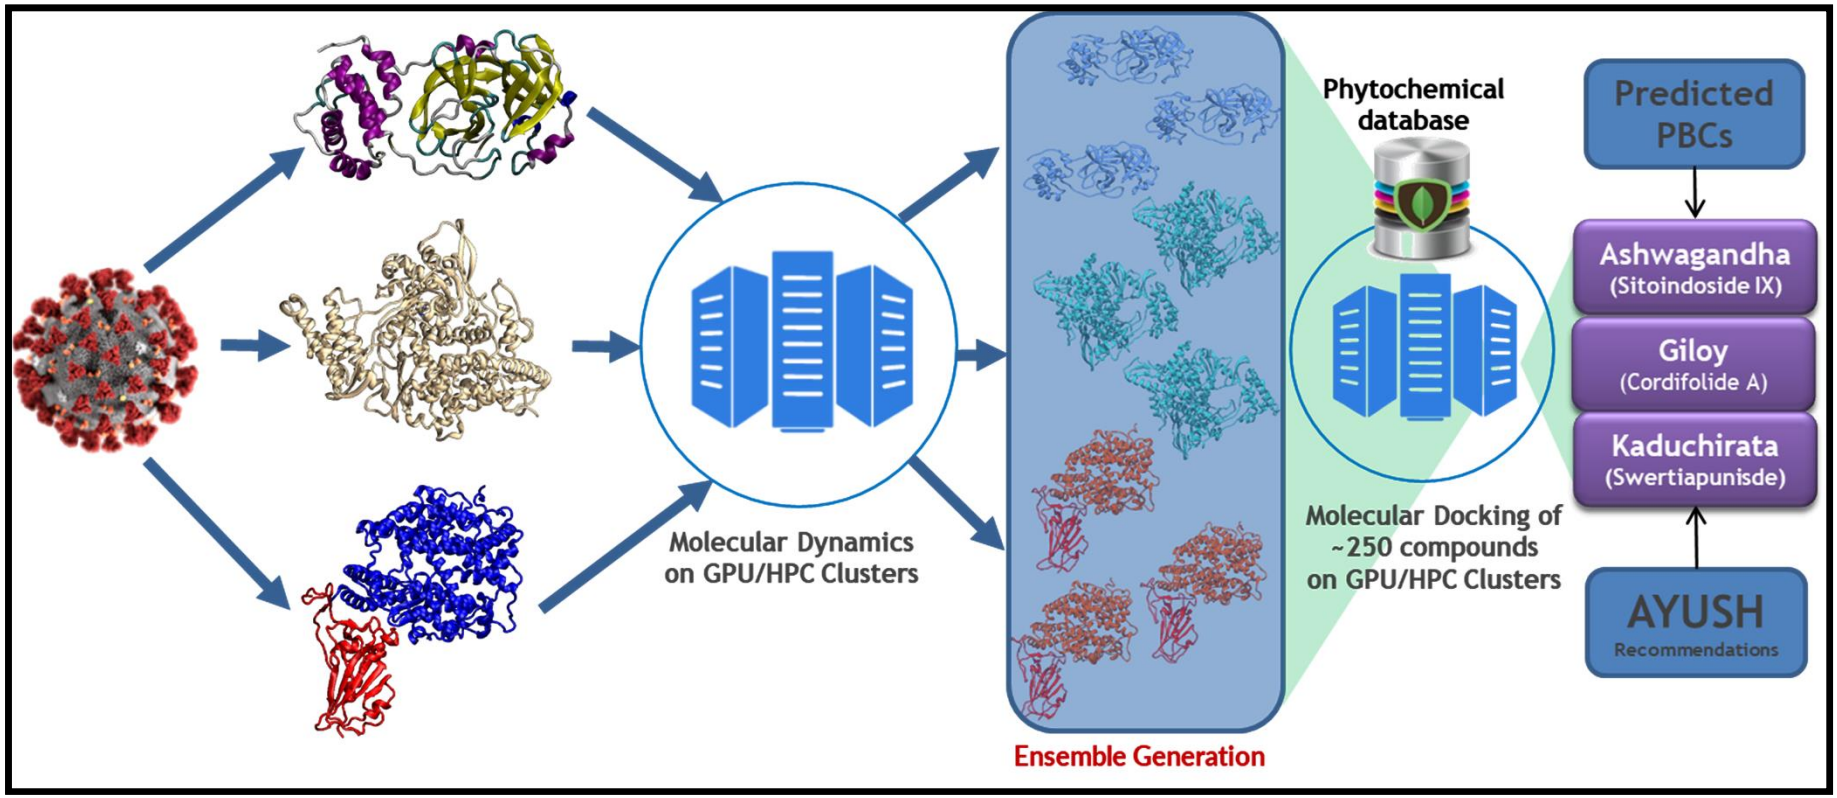

Thank you