

# TADaaM at INRIA Bordeaux

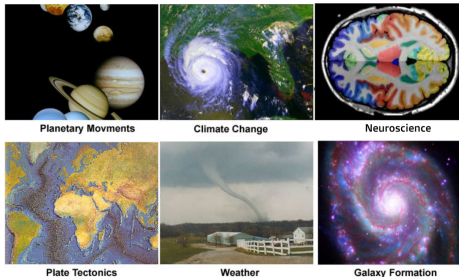
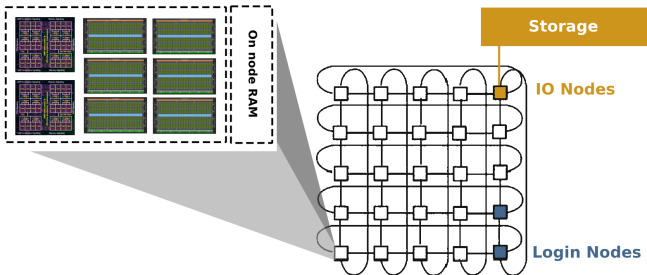
## Enabling the Deployment of Exploratory Applications on HPC Systems

Ana Gainaru

Vanderbilt University

## High-Performance Computing

- ▶ Large centralized machines
  - ▶ Used for large scientific applications ([https://en.wikipedia.org/wiki/Grand\\_Challenges](https://en.wikipedia.org/wiki/Grand_Challenges))
  - ▶ Capable of handling a lot of computation
- ▶ Specialized Software
  - ▶ Scheduler, Distributed Operating System, Filesystem



- ▶ HPC system evolve together with large monolithic applications
  - ▶ Focus on performance
  - ▶ Developed by the community for years
  - ▶ Tuned to scale and run on large-scale infrastructures



# Summit supercomputer



Located at Oak Ridge National Laboratory  
Currently the fastest supercomputer in the world

**200 PetaFlops peak performance** **4,608 nodes**: 2 IBM POWER9 CPUs and 6 Nvidia Volta V100 GPUs per node

- ▶ Each node contains two 22-core CPUs and 6 80-core coprocessors. **Total of 2,414,592 cores**
- ▶ Each node contains 608 GB of coherent memory
  - ▶ 800 GB of non-volatile RAM that can be used as a burst buffer or as extended memory
- ▶ **<1 - 5 days mean time between failures (MTBF)**



VANDERBILT  
UNIVERSITY

# Background



2011 - 2015



2015 - 2017



2017 - present

## PhD in Computer Science

Advisor: Marc Snir

Failure prediction, Hybrid checkpointing  
Fault tolerance framework for Blue Waters

## HPC Architect

Collective communication  
HW custom application optimization

## Research Assistant Professor

Scheduling  
Heterogeneous, dynamic applications

- ▶ 1 journal, 8 conference papers
- ▶ Part of SC13 panel on Fault Tolerance in HPC
- ▶ UIUC Academic Committee

- ▶ 1 conference paper and 1 patent
- ▶ PC member for 5 major conferences

- ▶ 2 journal, 2 workshop, 4 conference papers
- ▶ NSF Panel review member 2019/2020
- ▶ Editorial board for two HPC journals
- ▶ Vice-chair for SC Posters/Tutorials



VANDERBILT  
UNIVERSITY

# Research Statement



## HPC has created two new scientific paradigms

### Big data

- ▶ Dig through large amounts of data (sensors, transaction records, genome and protein databanks)

### Scientific computing

- ▶ Simulate physical phenomena when real experiments are very costly

## Motivated the development of major computational infrastructures

- ▶ HPC system evolve together with large monolithic codes
  - ▶ Application in Biochemistry, Chemistry, Physics, Environment Modeling
  - ▶ Each generation adding new technologies to increase the performance of applications
    - ▶ Accelerators, High-bandwidth memory, Burst buffers, GPU Direct, Unified Virtual Memory
  - ▶ Peak performance for system and applications has continued to increase
    - ▶ This comes with the price of higher degree of performance variation




VANDERBILT  
UNIVERSITY

# Research Statement

## Performance variability influences the effectiveness of the middleware

- ▶ At all levels of the software stack
  - ▶ Fault tolerance - During PhD
  - ▶ Data movement - During all my positions
  - ▶ Scheduling - During Vanderbilt



## ▶ New generation of HPC applications

- ▶ Exploratory applications from fields like neuroscience, bioinformatics, genome research, computational biology
- ▶ High variability in resource needs
  - ▶ within and between runs

**Research Goal: Bridging the gap between application needs and HPC systems**



## Follow the life cycle of an HPC application (and the corresponding data movement)

- ▶ Why data movement is important?

	Intrepid	Mira
Operational years	2008 - 2013	2013 -
Peak performance	0.6 PFlops	10 PFlops
Data transfer	90 GB/s	240 GB/s
Ratio between data transfer and Flops	160 GB/PFlop	17 GB/PFlop

1. The application is submitted on the HPC machine

2. During execution

- ▶ Fault tolerance
- ▶ I/O transfers
- ▶ Stochastic resource requirements

3. Application end: I/O dependencies



VANDERBILT  
UNIVERSITY

## 1. The application is submitted on the HPC machine

### ► Problems

#### ► 1.1. Applications need to request resources




Figure: Traces [2013-2016] of neuroscience apps (Vanderbilt's medical imaging database).

**Provide the optimal sequence of requests based on:**


(i) a model of the applications; (ii) a model of the platforms; (iii) resiliency schemes available



## 1. The application is submitted on the HPC machine

### ► Problems

- 1.1. Applications need to request resources
  - The model is in the testing phase by the neuroscience departments in Vanderbilt
- 1.2. Pre-fetch data before application start



### Directions:

- Using ORNL ADIOS I/O middleware to define the I/O dependencies
- Using the model to decide which HPC resource will become available
- Speculatively move the data to the most likely location
  - Trade-off between creating more data movement and having the data ready



VANDERBILT  
UNIVERSITY

## 2. During execution


### Problems

#### ► 2.1. Time sensitive applications

- Applications that require the result in real time
- Data is gathered by sensors that can be geographically distributed from the compute nodes
- **Direction**
  - Develop I/O middleware capable of prioritizing data transfers
  - Overlap transfer and computation based on application properties

#### ► 2.2. Stochastic applications

- Applications that do not need the same resources throughout their execution



- Memory footprint, processing units
  - E.g. Simulations that use adaptive mesh refinement to refine the accuracy of their solutions could require an additional 128 CPU cores during its computation to increase the resolution in an area of interest
- Using adaptive MPI and models to define co-scheduling policies




## 2. During execution

### ► Problems

#### ► 2.3. Fault tolerance

	MTBF	Time to save the state of an application
Today's systems	hours - few days	few hours
Prediction for future systems	minutes	minutes

It is expected that the time to take a checkpoint will exceed the MTBF



Failure type	Percentage	Crashes	Degradation
Luster MDT failure	39.6%	5%	2.7x
Luster OST failure	16.3%	13%	3.3x
DIMM failures	15.7%	11%	1.4x

► Failure prediction to decide when and where to take checkpoints

► Integrated in the Blue Waters software stack

► Failures that cause performance degradation




VANDERBILT  
UNIVERSITY

## 2. During execution

### Problems

#### ► 2.4. I/O congestion

- Pre-fetching, taking snapshots of applications increase the data movement in the system
- Shared communication and I/O networks
- **Direction**
  - Schedule the I/O of applications
  - Optimize the use of hardware for data movement



- Online as applications progress
  - Need a centralized control unit to monitor traffic
- Within the batch scheduler based on past behavior of applications

#### ► Stochastic applications

- Need to adapt to shifts in behavior



## Research Goal: Bridging the gap between application needs and HPC systems

- ▶ **Long term:** Develop I/O middleware for optimizing the execution of HPC applications in the presence of high performance variability at all levels of the software stack

## Method

- ▶ **Understand performance variability**
  - ▶ Intrinsic to the application (resource requirements)
  - ▶ I/O congestion (connected with fault tolerance and scheduling)
- ▶ **Communication between the application and the middleware**
  - ▶ Application level fault tolerance to minimize I/O transfers
  - ▶ Data pre-fetching
- ▶ **Design new middleware to adapt to the needs of HPC applications**
  - ▶ Design a new paradigm for data transfers for time sensitive applications
    - ▶ Choose the granularity of the transfers (trade-off between result precision and time)
    - ▶ Adapt compression schemes
  - ▶ Adapt current middleware for high performance variability
    - ▶ Adaptive / speculative executions
    - ▶ Application dependencies

Leverage TADaaM expertise with IO Scheduling, MPI, hwloc



VANDERBILT  
UNIVERSITY

## ► PhD Thesis

- Failure prediction for HPC systems
- Research Assistsanship with NCSA (my work was used by the Blue Waters system, 2012 - 2019)
- Panel of experts at SC13

## ► Service

- UIUC Academic Integrity Appeals Committee, Engineering Graduate Student Advisory Committee
- NCSA Colloquium Committee
- Student mentoring (Vanderbilt)

## ► Post-doc research

- Performance optimization for scientific applications
- Tools developed used by neuroscience applications
- Mellanox / Vanderbilt

## ► Teaching

- Politehnica University of Bucharest / UIUC / Vanderbilt
- OS, Parallel and distributed computing, HPC

## ► Editorial and review boards

- NSF Panel review 2019/2020
- Editorial borad for the international journals IJHPCA, JPDC, TPDS
- PC member for 15+ conferences

## ► Chairs

- Co-chair for Posters SC18, Tutorials SC20, FTS 2017 Workshop
- Organizer for CCIW Workshop 2019

## ► Publications

- 18 conference, 5 journal, 2 workshop
- 1 patent (Mellanox)

## ► Research Statement

- I/O middleware for optimizing the execution of HPC applications in the presence of high performance variability at all levels of the software stack



VANDERBILT  
UNIVERSITY