Michael Lewis

PhD Candidate

Disasterville

First person educational game designed for middle-scholars to understand what to do and not to do in the home when there is an earthquake.
Funded by the Illinois Terrorism Task force.

Distributed Computing Research

Distributed Classes	Systems/Frameworks	Key features
Grid	Condor, Globus	Ability to find resources. Virtualization technology.
Cloud	Open Stack, AWS, Azure, Hadoop, Spark	Transparency of communication and I/O, Persistence of data.
HPC	MPI, UPC	Developer specification of communication and I/O

Open Science Data Cloud

2014 – Sao Paulo Brazil

2014 – Amsterdam, Netherlands

2013 – Sao Paulo Brazil

2013 Edinburgh Scotland

Resource Description Framework (RDF)

RDF dataset is a collection of triples.

Triples are Subject-Predicate-Object entities.

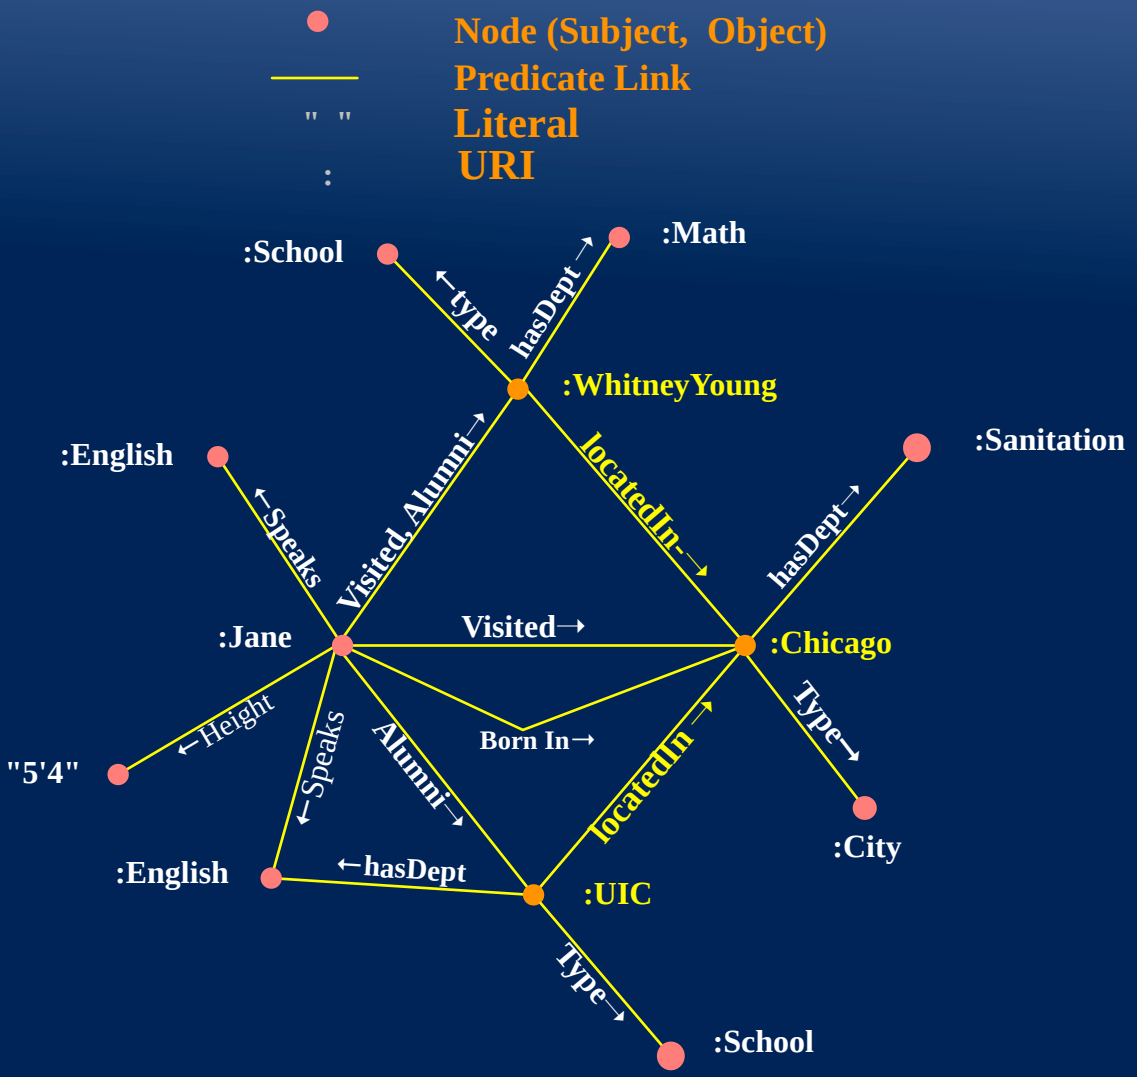

Triple Patterns

Classification Patterns

UIC type Unversity

Illinois type state

Mike type person

Pattern 3

Mike Born in Chicago

Pattern 1

UIC current-student Mike

Pattern 2

UIC located in Illinois

Path approach to processing and retrieving linked data within a distributed environment.

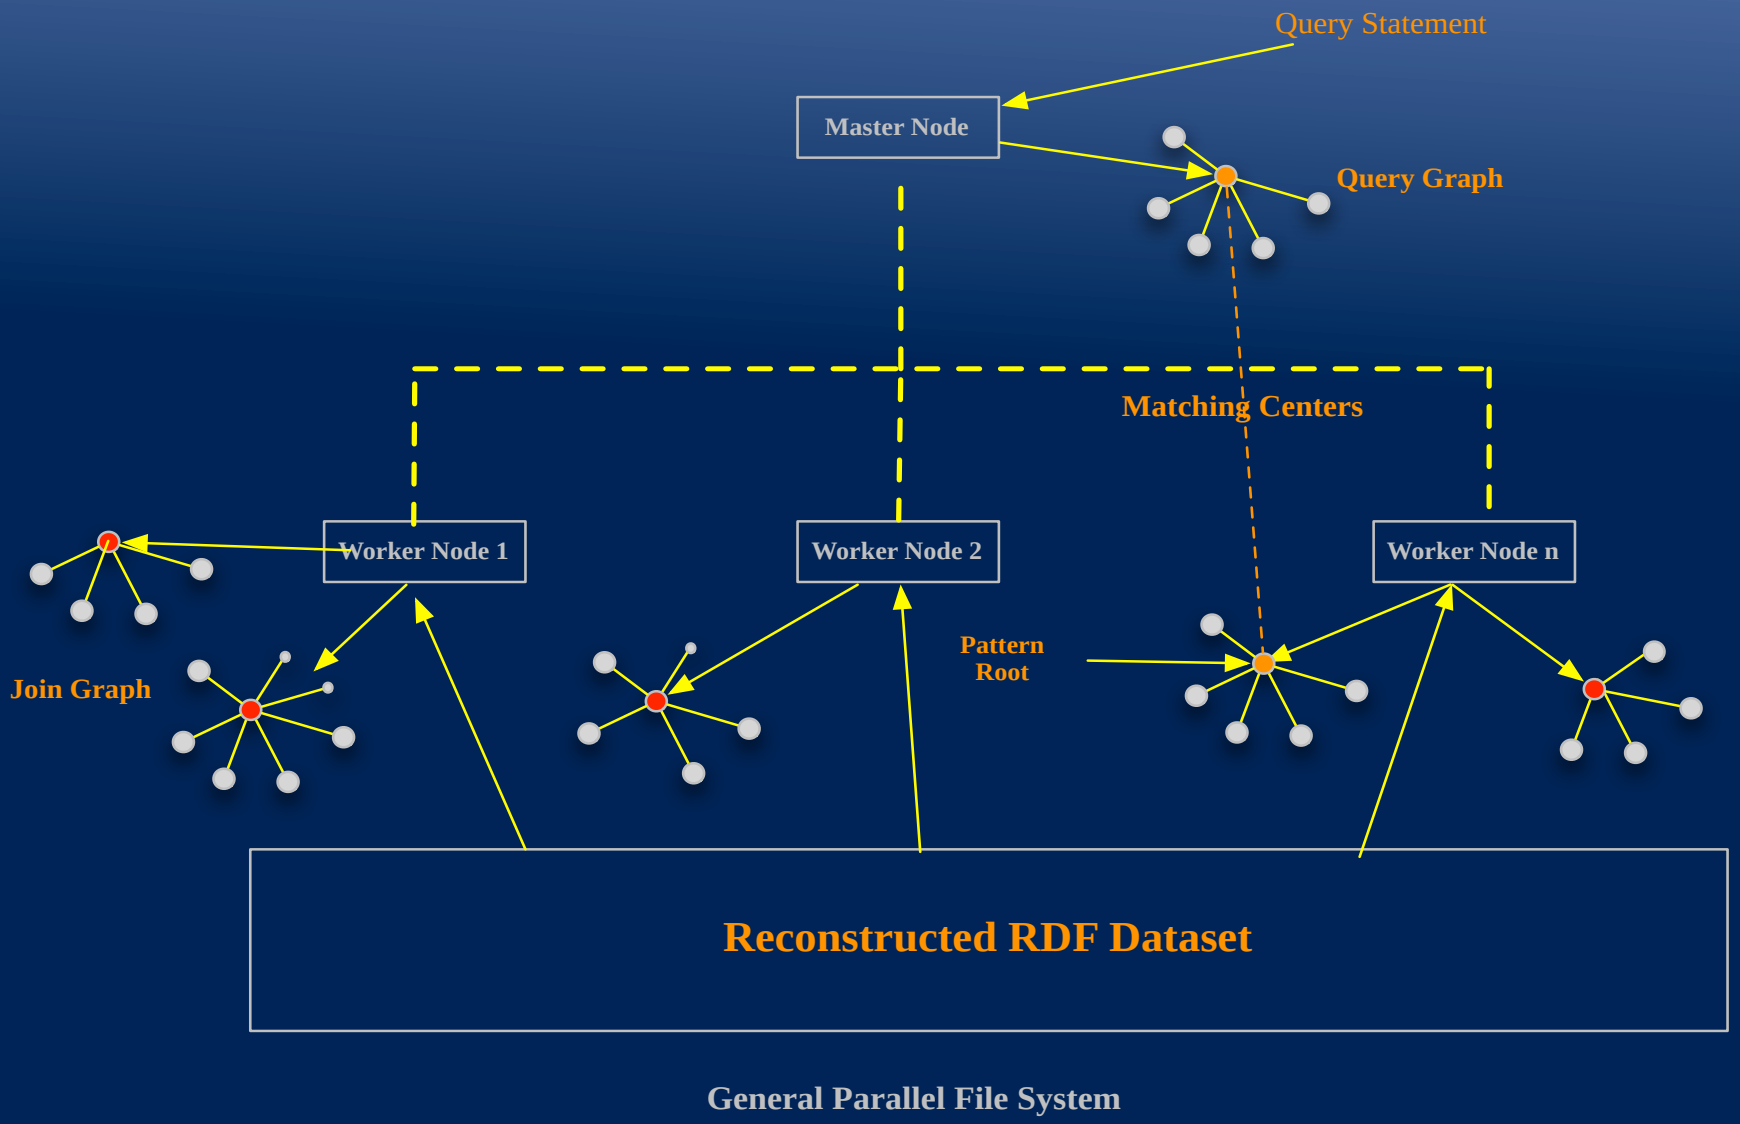

Argonne Internship

Advisor: Preeti Malakar

Research Area: Optimizing Read Write times for Multidimensional Datasets

Architecture

- Master-Worker Architecture. Using MPI
- Cooley
  - 126 Compute Nodes
  - 2 Intel (2.4 Ghz processors) per node. 6 Cores per CPU
  - 24 GB Ram per node 12 GB per GPU.
- Mira – an IBM Blue Gene Supercomputer
  - 49, 152 compute nodes.
  - 16 Cores per node, at 1.6 GHz.
  - 16 GB of DDR3 memory.
- General Parallel File System (GPFS)
  - 24 PB of capacity.
  - 240 GB/s bandwidth.
  - Allows data to be accessed over a multiple number of computers at once.

Data Decomposition in S3D

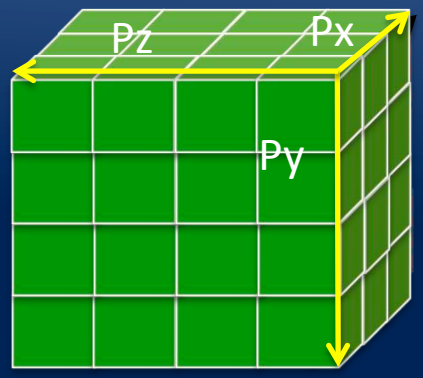

Data Volume

S3D
- Combustion simulation code
- Uses 3D data volume for simulation

Data Decomposition : Nx x Ny x Nz (E.g. 12 x 12 x 12)

Virtual Process topology : P_x x P_y x P_z (E.g. 3 x 4 x 4)

Default I/O mechanism in S3D:
Each process performs I/O for its corresponding sub-blocks of data

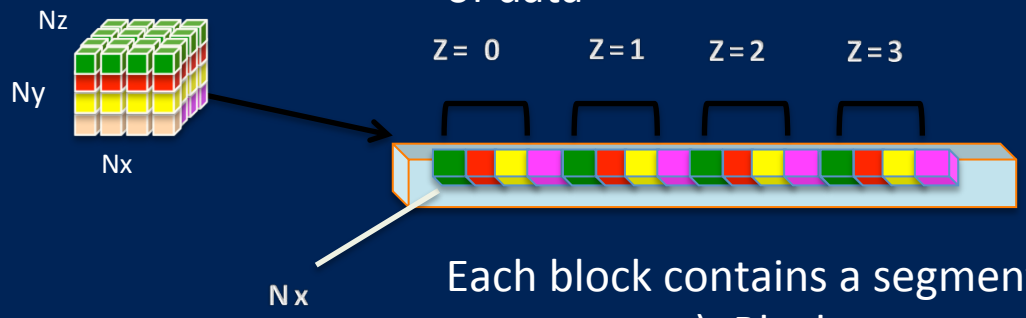

Each block contains a segment of solution data (E.g. pressure, temperature). Blocks are stored contiguously in the filesystem.

I/O through few processes

Hiero / Fixed Plane (Previous work): Selects X=0 processes for I/O
E.g.: 24 processes (leaders) from x = 0 plane for reading its entire row data.

Our approach: Extend HieRO to optimally select the leaders
Smallest Plane: Selects smallest plane (fewest processes)
E.g.: Selects 12 processes from the z = 0 plane for reading.

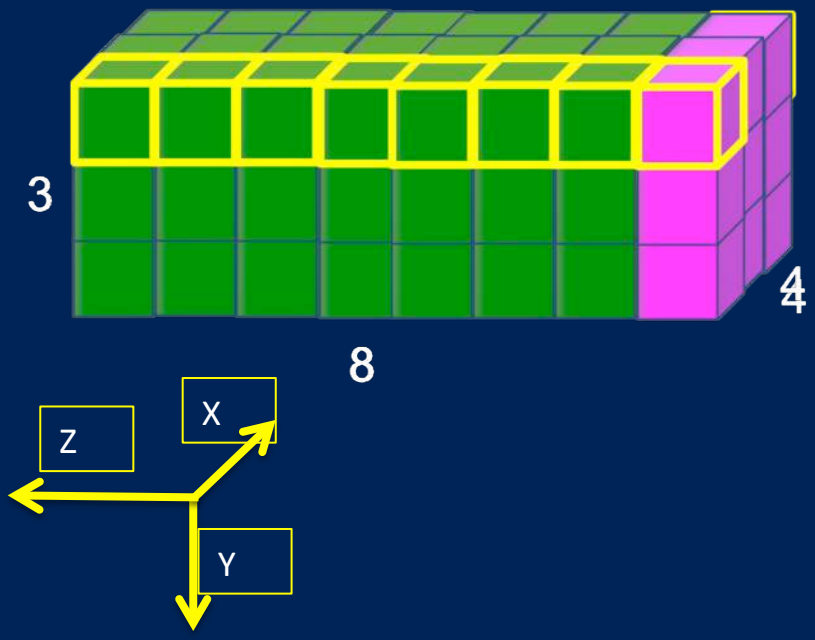

Why is this done ?
- further reduce contention

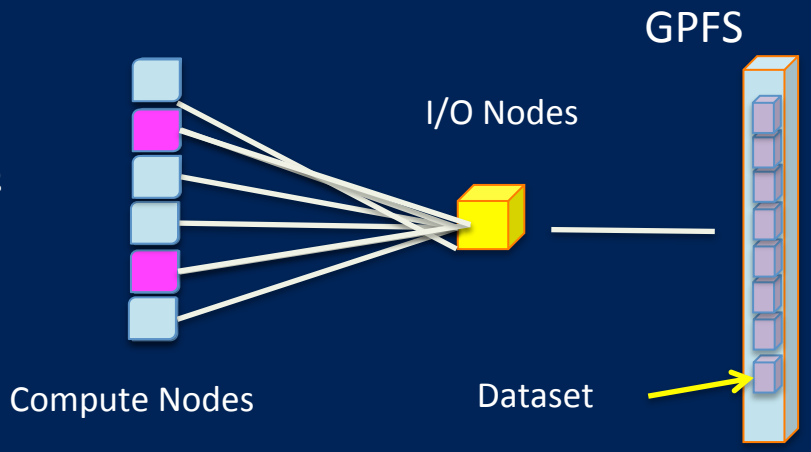

Thank You !!