QUERY-DRIVEN SPATIOTEMPORAL DATA ANALYTICS

Nagiza F. Samatova
North Carolina State University
Oak Ridge National Laboratory
samatova@csc.ncsu.edu

NSF Expeditions, AHM@UMN
May 26-27, 2011

OAK RIDGE National Laboratory

NC STATE UNIVERSITY
Department of Computer Science

Publications Covered in This Talk

- ***Srisam Lakshminarasimhan***, ***Neil Shah***, S. Ethier, S. Klasky, R. Latham, R. Ross and N. F. Samatova, *Compressing the Incompressible with **ISABELA**: In-situ Reduction of Spatio-Temporal Data,* 17th International European Conference on Parallel and Distributed Computing [Euro-Par 2011], **Nominated for Distinguished Paper Award**.

- ***Sriram Lakshminarasimhan***, ***John Jenkins***, ***Isha Arkatkar***, Z. Gong, H. Kolla, S-H Ku, S. Ethier, J. Chen, CS Chang, S. Klasky, R. Latham, R. Ross and N. F. Samatova, ***ISABELA-QA**: Query-driven Data Analytics over ISABELA-compressed Scientific Data*, [Supercomputing 2011], **Under Review**.

- **List of all the other 10 publications in the NSF report of 5/25/11**

Publications Covered in This Talk

- ***S. Lakshminarasimhan***, ***N. Shah***, S. Ethier, S. Klasky, R. Latham, R. Ross and N. F. Samatova, *Compressing the Incompressible with **ISABELA**: In-situ Reduction of Spatio-Temporal Data,* 17th International European Conference on Parallel and Distributed Computing [Euro-Par 2011], **Nominated for Distinguished Paper Award**.

- ***S. Lakshminarasimhan***, ***J. Jenkins***, I. Arkatkar, Z. Gong, H. Kolla, S-H Ku, S. Ethier, J. Chen, CS Chang, S. Klasky, R. Latham, R. Ross and N. F. Samatova, ***ISABELA-QA**: Query-driven Data Analytics over ISABELA-compressed Scientific Data*, [Supercomputing 2011], **Under Review**.

Motivating Factors

- Current large scale scientific simulation can easily produce several terabytes during a single run.

- I/O accesses have become a rate limiting factor to key scientific discoveries.
  - A paradigm shift is necessary for analysis in exascale.

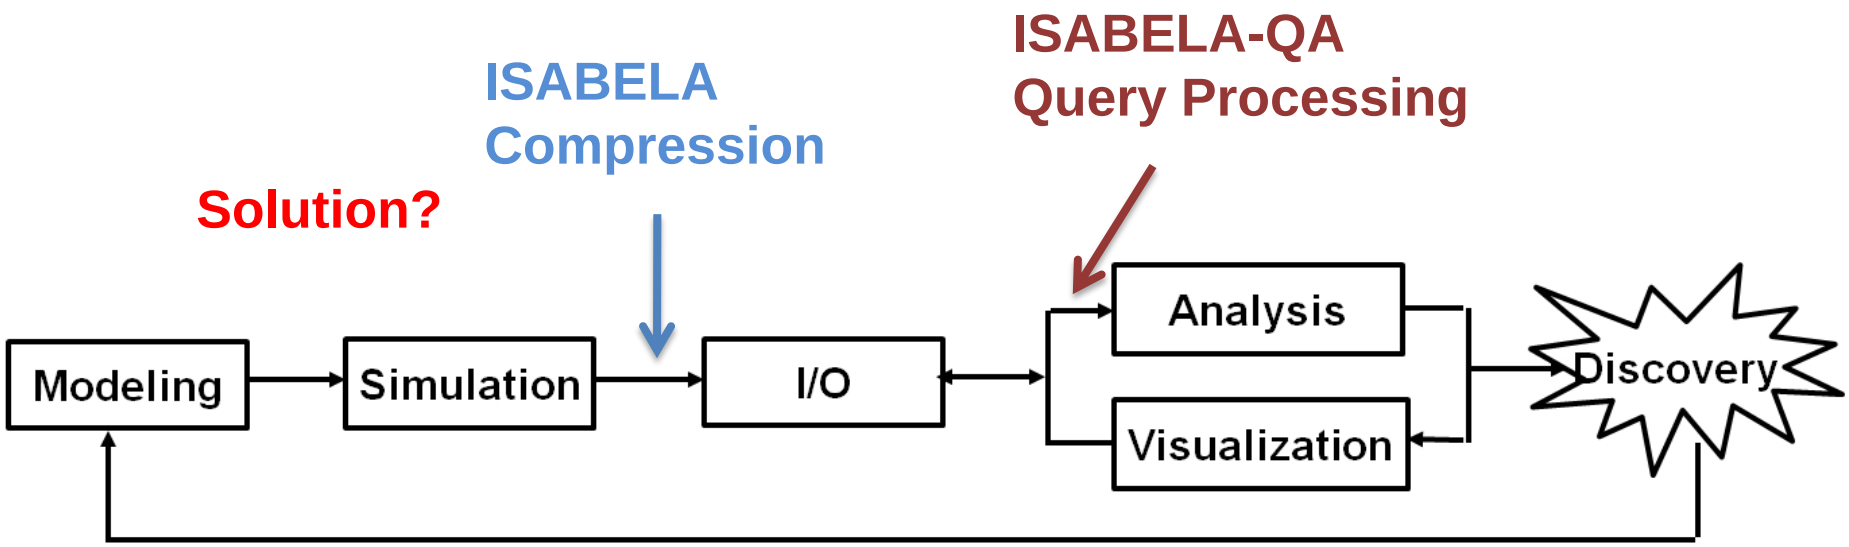

Typical Simulation Data

Expected: 10-fold increase by 2012-2014

C&R Data

- ~2TB per Checkpoint & Restart
- Every 1 hour
- Two copies
- Keep the last copy
- Accessed once or never.

Lossless

Analysis Data

- ~TB per run (now)
- Every 10th or 100th time step
- Repeated access.
- Analysis routines and I/O reads
- Matlab analysis routines

Inherently Lossy

V&V Data

- ~ GBs
- Every 2nd time step
- Verify, and validate simulation.

Myth: "Scientific data is almost uncompressible" with lossless encoding

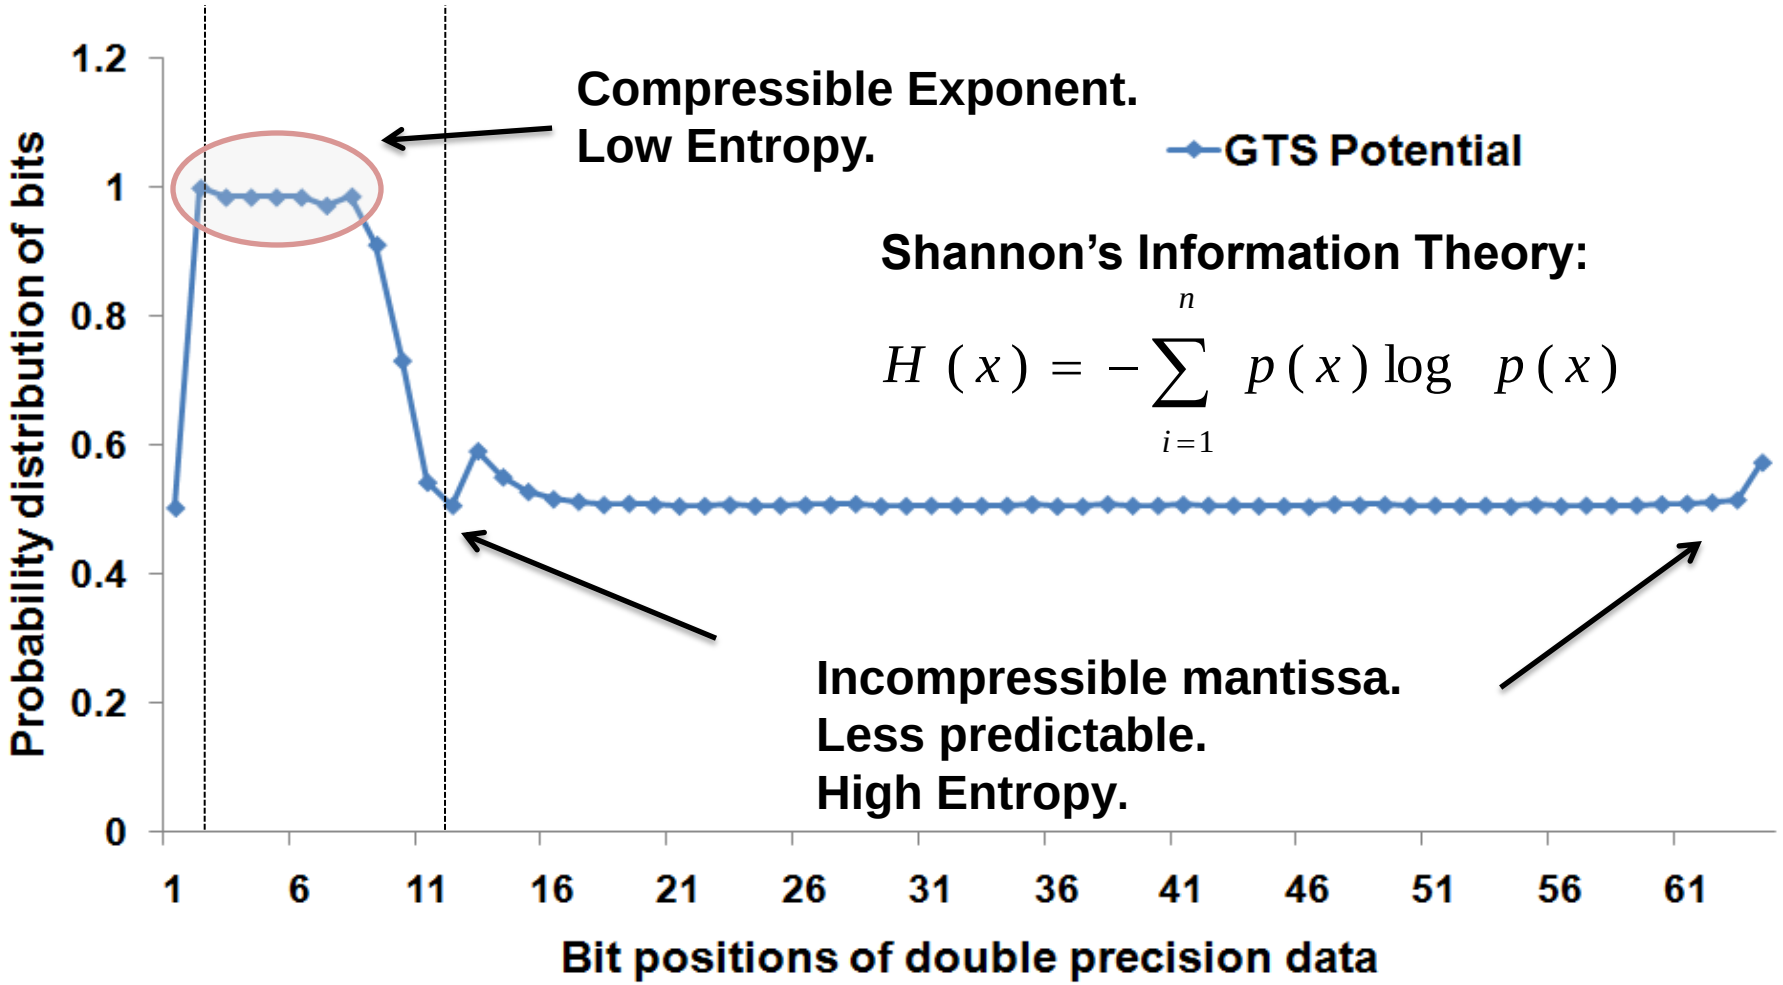

* Courtesy: Eric Schendel

Still “Incompressible” with Lossy Encoding

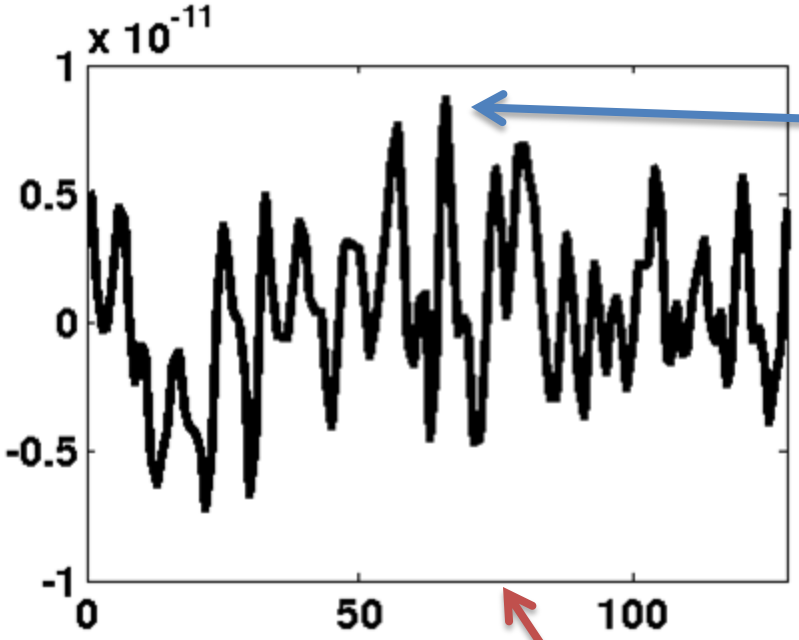

Original Data

- Highly random

- Extreme events missed

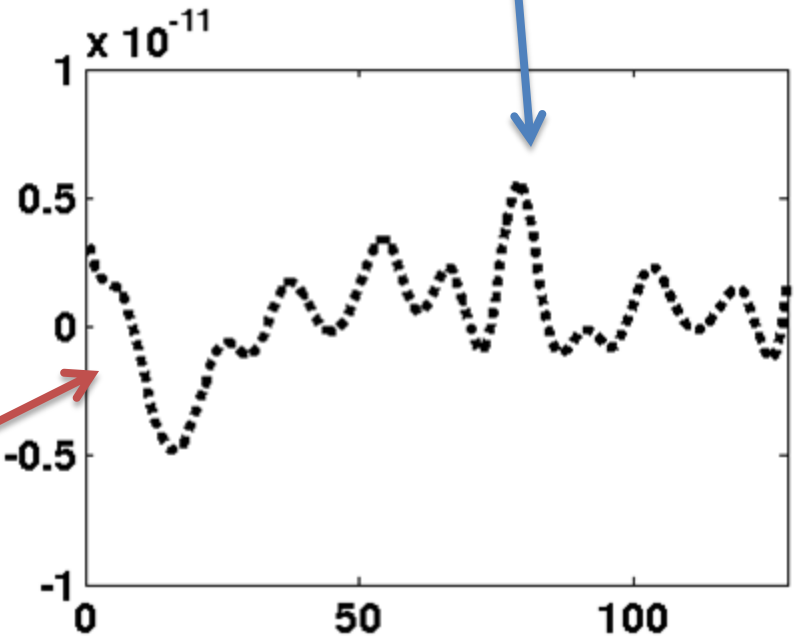

B-Spline/Wavelet Reduction

~0.4 correlation!

Compression Techniques Performance Comparison

- GTS dataset, 170K vector, potential.

- Fixed overall correlation (>0.99*) & NRMSE (<0.01*) for lossy techniques

Metric	FPC	LZMA	ZIP	BZIP2	Wavelets*	B-splines*	ISABELA*
Lossless	Yes	Yes	Yes	Yes	No	No	**No**
CR_M	3.71	2.72	1.13	1.11	22.51	0	**81.44**
Compression (sec.)	0.58	7.01	1.03	3.96	0.62	0.78	**0.93**
Decompression (sec.)	0.56	1.38	0.49	1.18	0.58	0.82	**1.05**

ISABELA's Magic Bullet

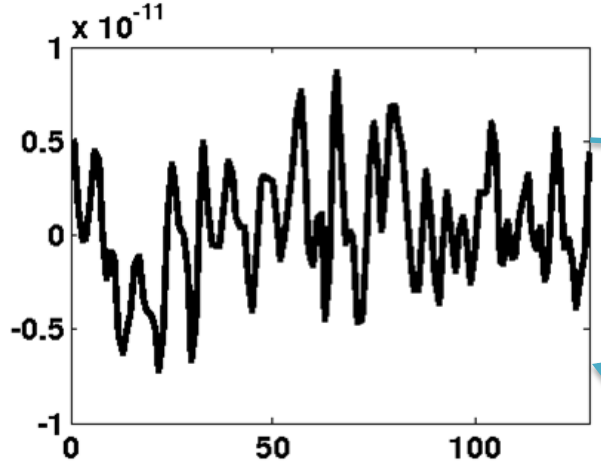

Sort Preconditioner

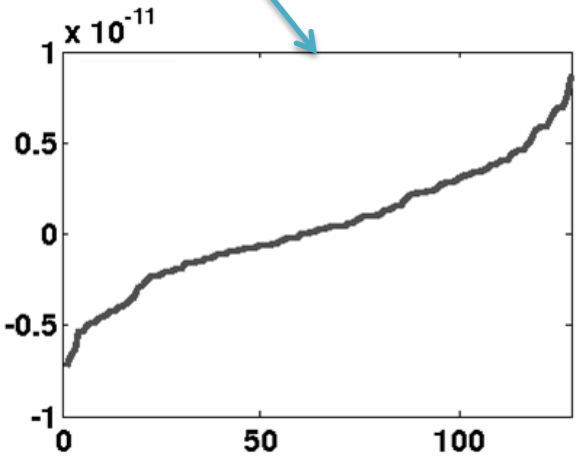

~0.99 correlation!
0.001 NRMSE

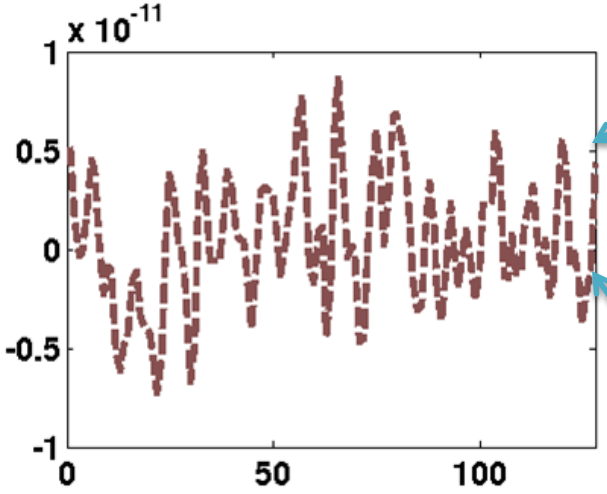

B-spline Reduction

OAK RIDGE National Laboratory

NC STATE UNIVERSITY Department of Computer Science

ISABELA: Communication-Free Workflow

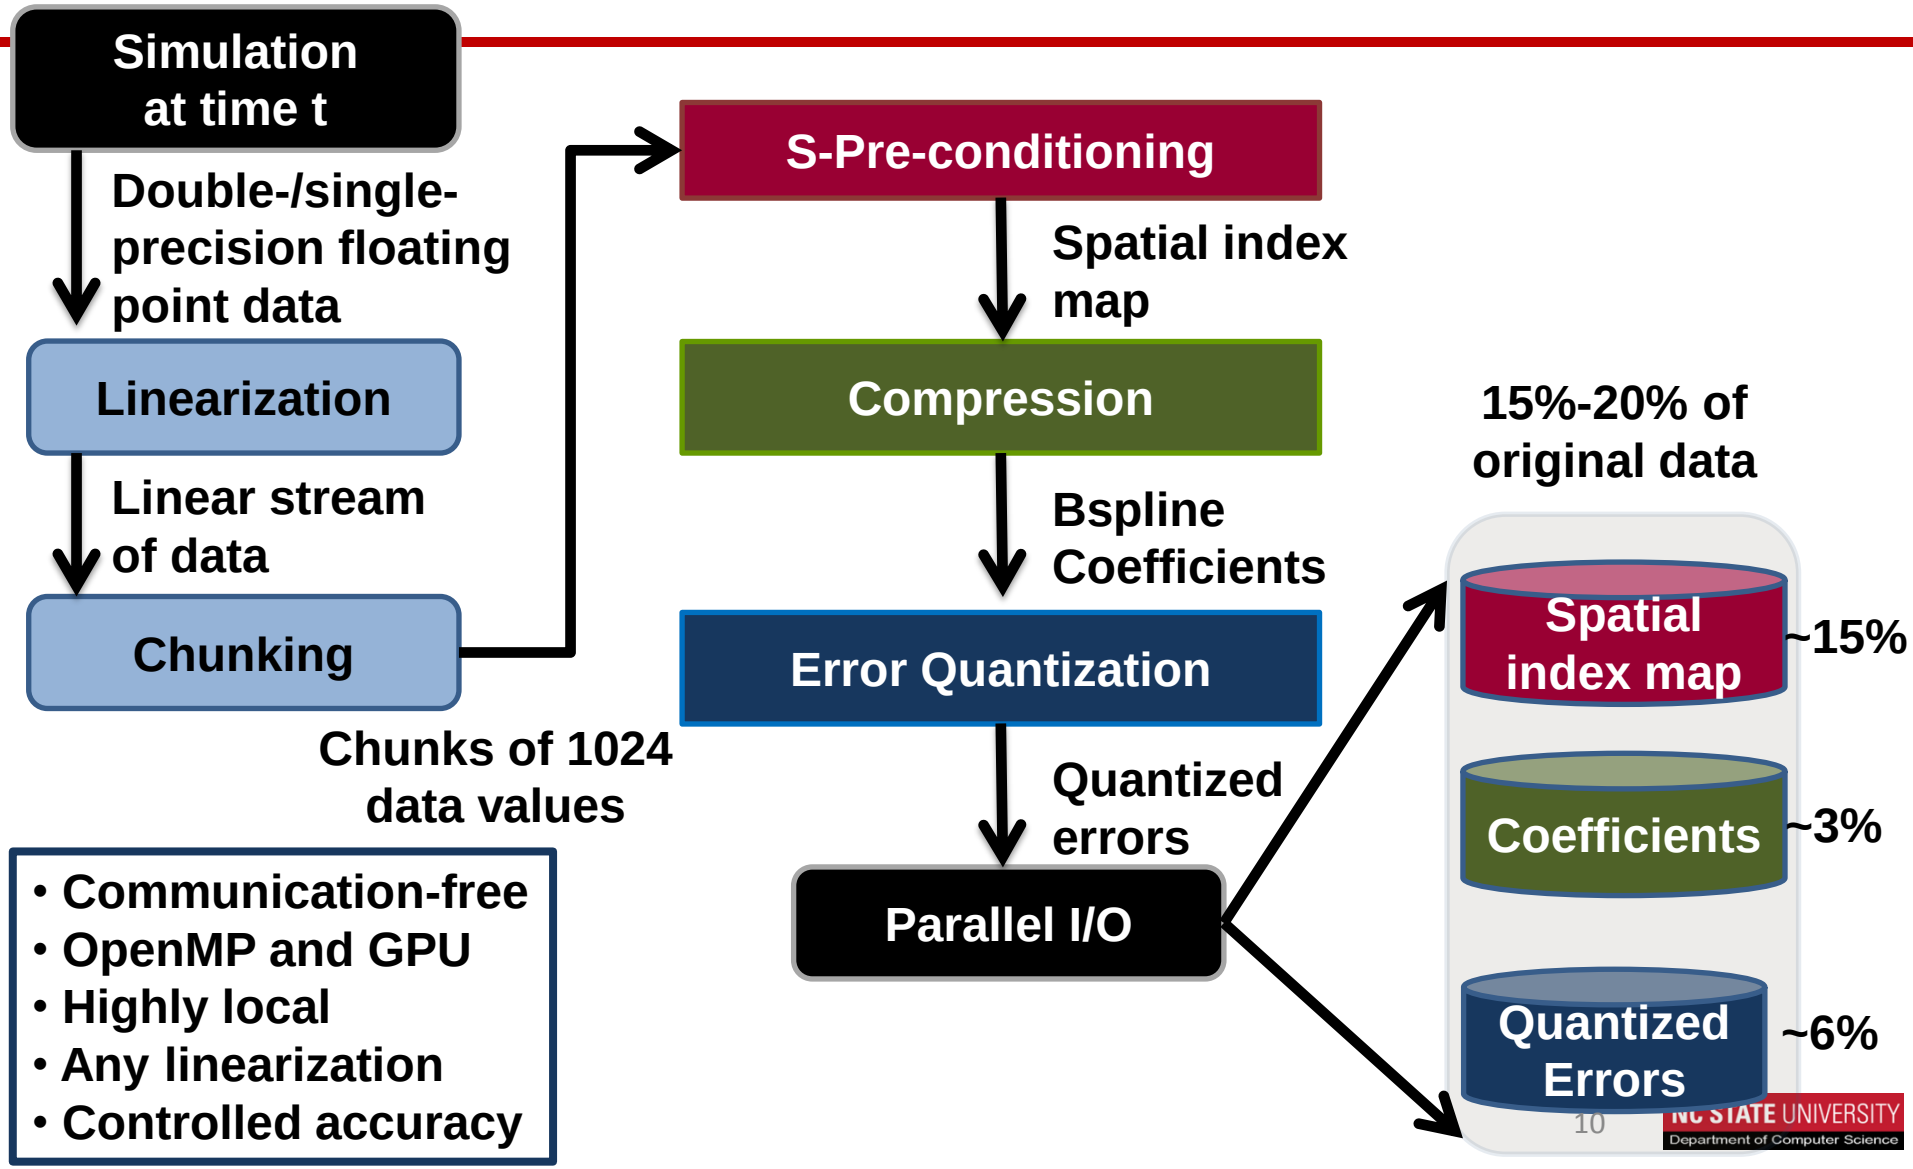

Accuracy Robustness (CR=80%)

- **ISABELA consistently gives almost a near-perfect correlation in Linear stages of the GTS simulation.**

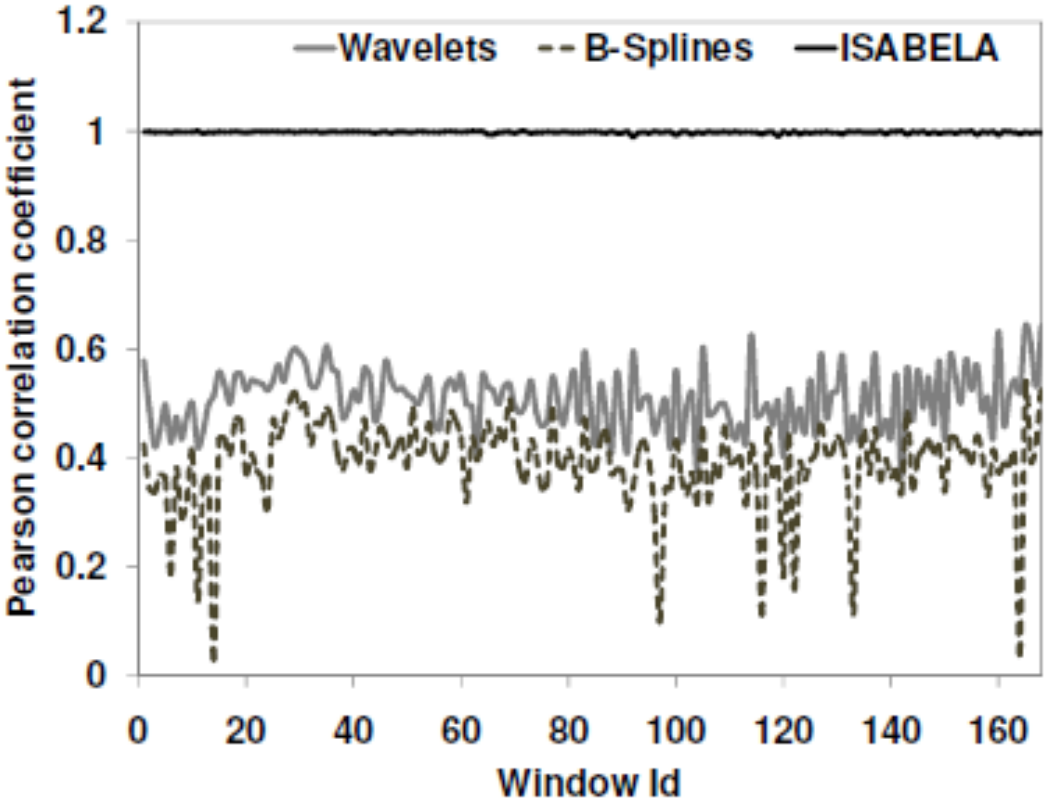

Performance on Public Datasets (CR=80%)

	Correlation				NRMSE			
	WAVELETS		ISABELA		WAVELETS		ISABELA	
	Average	±	Average	±	Average	±	Average	±
msg_sppm	0.4	0.287	**0.982**	**0.017**	0.203	0.142	**0.051**	**0.015**
msg_bt	0.754	0.371	**0.981**	**0.054**	0.112	0.151	**0.038**	**0.024**
msg_lu	0.079	0.187	**0.985**	**0.031**	0.422	0.103	**0.048**	**0.015**
msg_sp	0.392	0.44	**0.967**	**0.051**	0.307	0.243	**0.064**	**0.033**
msg_sweep3d	0.952	0.07	**0.998**	**0.006**	0.075	0.036	**0.004**	**0.003**
num_brain	**0.994**	**0.008**	0.983	0.028	**0.01**	**0.011**	0.011	0.005
num_comet	0.988	0.018	**0.994**	**0.025**	0.02	0.02	**0.01**	**0.006**
num_control	0.614	0.219	**0.993**	**0.017**	0.083	0.037	**0.009**	**0.002**
num_plasma	0.605	0.062	**0.994**	**0.004**	0.277	0.038	**0.033**	**0.004**
obs_error	0.278	0.203	**0.994**	**0.004**	0.303	0.091	**0.024**	**0.009**
obs_info	0.717	0.136	**0.993**	**0.006**	0.181	0.078	**0.026**	**0.016**
obs_spitzer	**0.992**	**0.001**	0.742	0.004	**0.005**	**0.001**	0.03	0.032
obs_temp	0.611	0.114	**0.994**	**0.011**	0.096	0.025	**0.009**	**0.003**

Effect of ISABELA Compression on Data Analysis with Combinations of Decompressed Variables

Analysis of XGC1 Potential Data

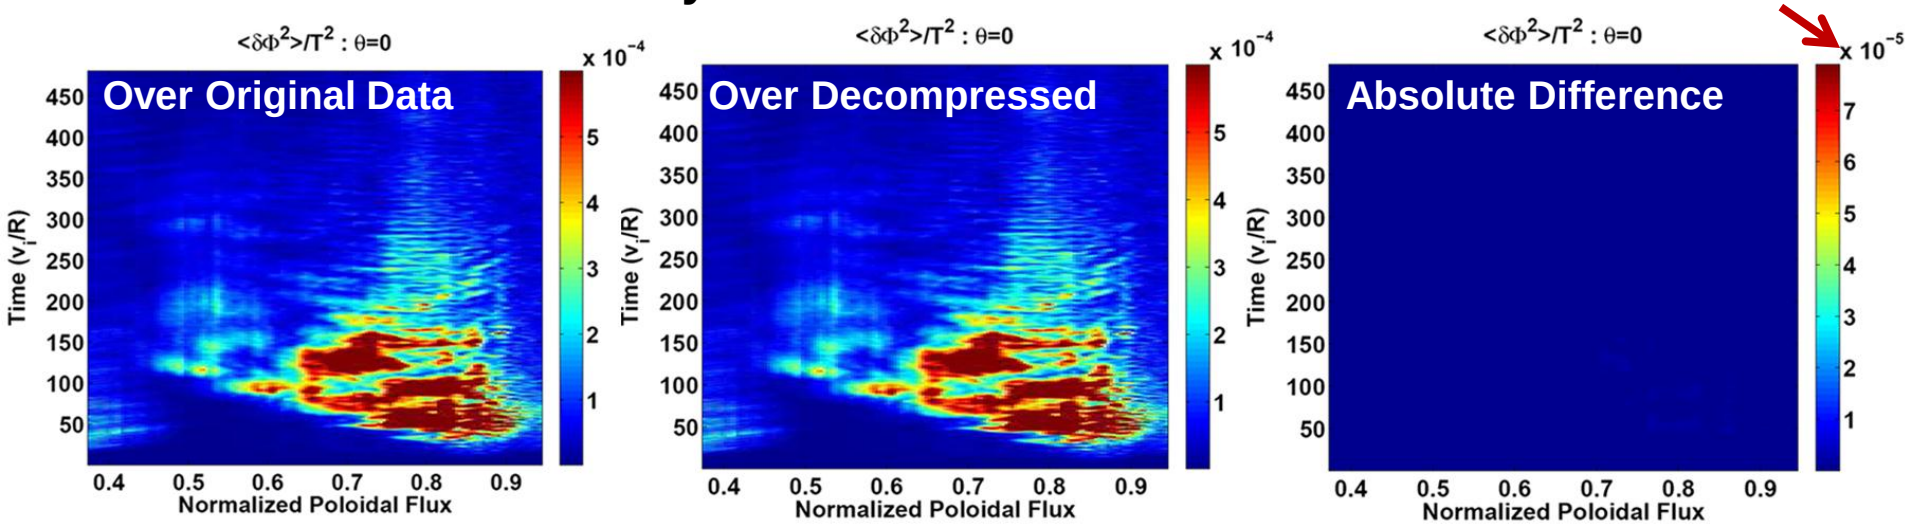

Analysis of S3D Temperature over 20M Grid Points

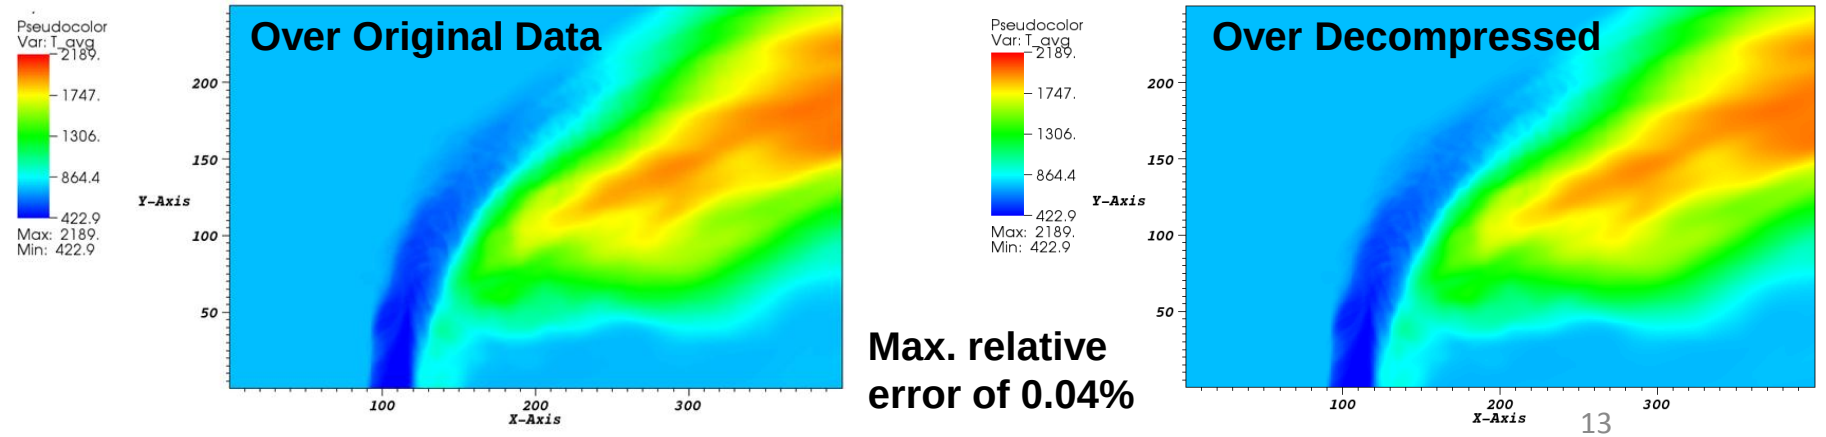

Max. relative error of 0.04%

National Laboratory

*** Courtesy Hemanth Kolla, Jackie Chen @ SNL, C.S. Chang, S-H. Ku @ NYU**

Publications Covered in This Talk

- ***S. Lakshminarasimhan***, N. Shah, S. Ethier, S. Klasky, R. Latham, R. Ross and N. F. Samatova, *Compressing the Incompressible with **ISABELA**: In-situ Reduction of Spatio-Temporal Data,* 17th International European Conference on Parallel and Distributed Computing [Euro-Par 2011], **Nominated for Distinguished Paper Award**.

- ***S. Lakshminarasimhan***, ***J. Jenkins***, ***I. Arkatkar***, Z. Gong, H. Kolla, S-H Ku, S. Ethier, J. Chen, CS Chang, S. Klasky, R. Latham, R. Ross and N. F. Samatova, ***ISABELA-QA**: Query-driven Data Analytics over ISABELA-compressed Scientific Data*, [Supercomputing 2011], **Under Review**.

Motivation for Query-driven Data Analytics

- **Exploratory data analytics is global-context and thus I/O bound:**
  - Need *full context* of simulation data, not just a single time step.
  - Data-intensive analytics is constrained by I/O.

- **Promise: Query-driven data reduction for downstream analysis**

- **Domain scientists' knowledge priors:**
  - **Climate**: Find Correlations between sea surface temperature anomalies and hurricane activity in **North Atlantic region**
  - **Fusion**: Calculate the turbulence intensity flux at **each radial zone** and their time correlations
  - **Fusion**: Calculate the time correlations of turbulence intensities **between different radial zones**.

$$0.1 < \psi < 0.15$$

$$0.3 < \psi < 0.35$$

$$0.5 < \psi < 0.55$$

$$0.7 < \psi < 0.75$$

$$0.9 < \psi < 0.95$$

Challenges with B-Trees

- DBMS/ Indexing based techniques make it easy to **locate** "needles in a haystack"

- But analysis still remain difficult.

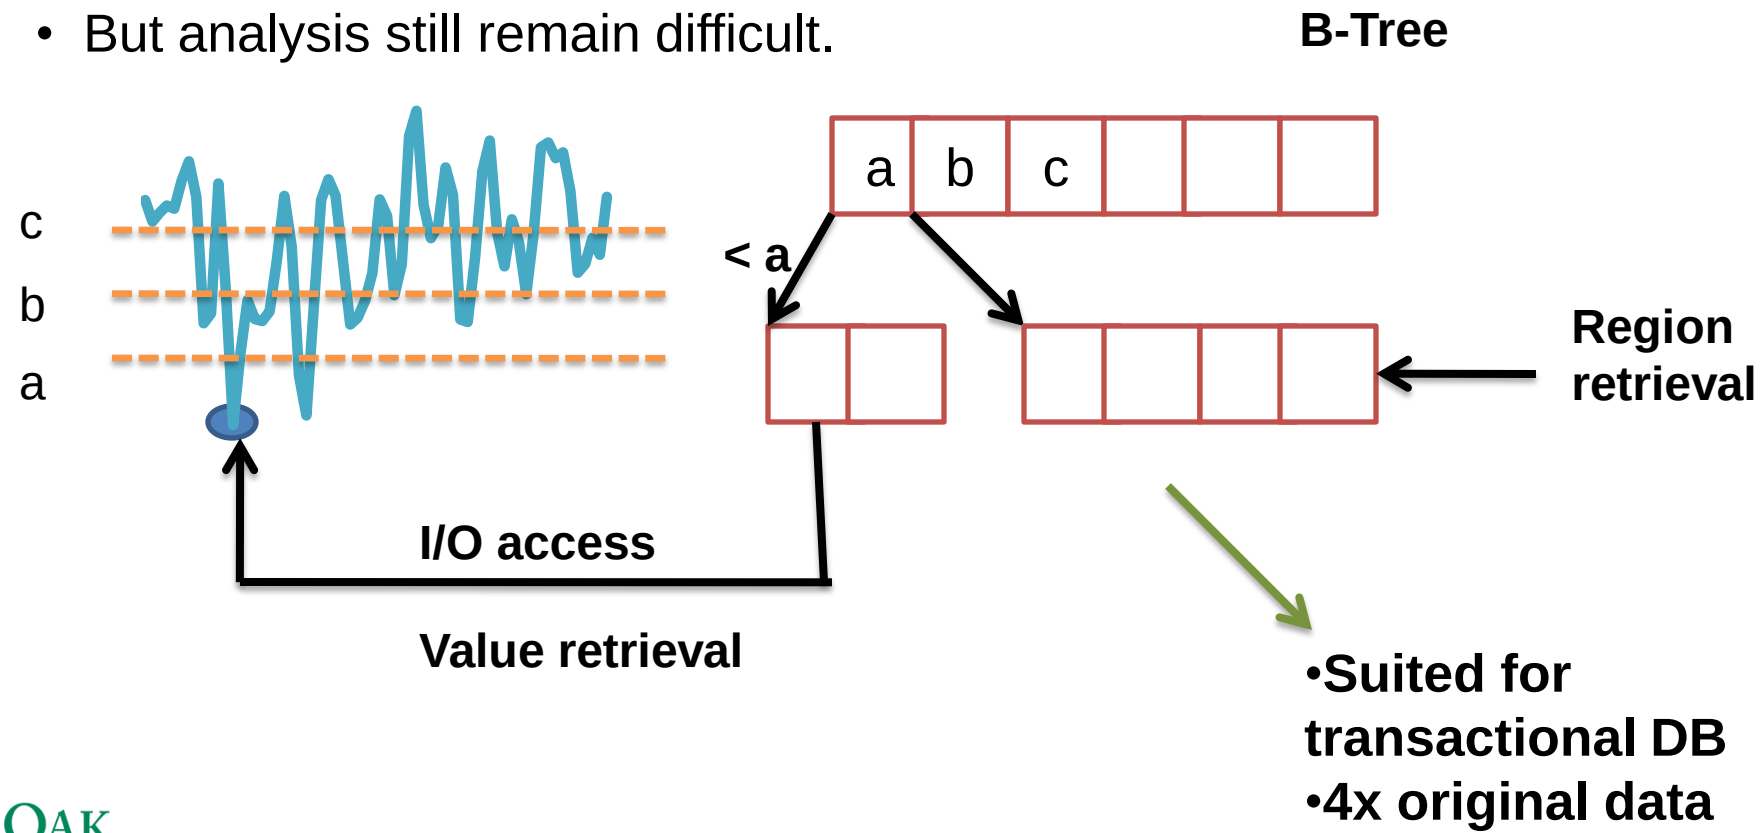

Challenges with Fastbit

Record # / Bin	(-∞, a)	[a, b)	[b, c)
1	0	1	0
2	0	1	0
3	0	1	0
4	0	1	0
5	1	0	0
6	1	0	0
7	0	0	1

- Tailor made for read-only databases.
- Bitmap-based indexing.
- Encode each bitmap vector
  - ***WAH*** compressed
  - Ability to perform fast logical operations directly on compressed space.
- Index size is **0.3 – 1.7x**
- **I/O** access dominate query cost

Challenges for Query-Based Analytics

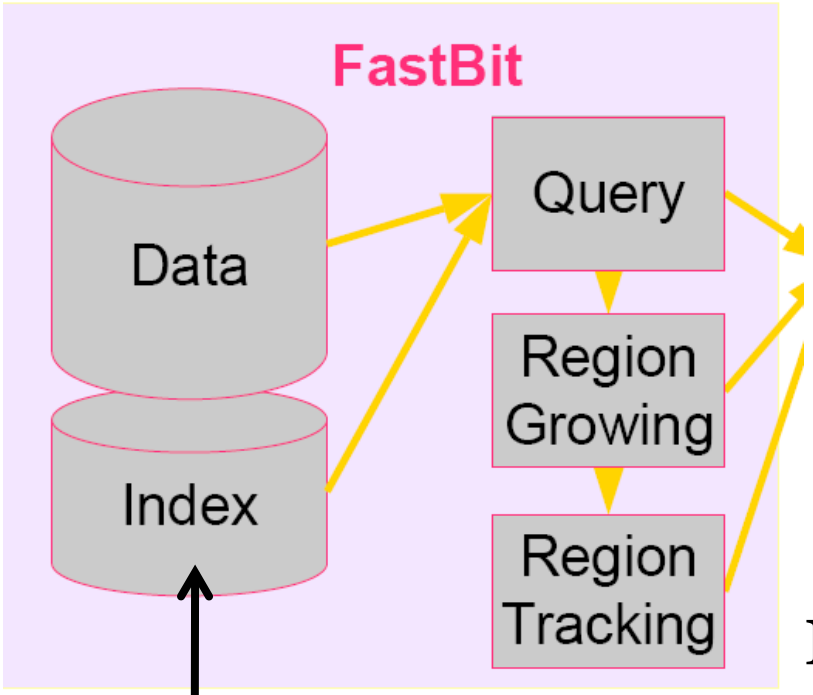

33%—170% of Original Data

Multiple Variables over Different Spatial Resolution

INDEX over Common Spatial Resolution

Query

Regions
Fast
Specific Analysis (e.g., histogram)

Quantized Data
Which analyses are robust to quantization?

Raw Data
I/O-bound

LIMITATIONS OF EXISTING INDEXING APPROACHES

- **Optimized for region retrieval**
  - Output: points in space that meet the query (50<Temperature<70 & 250<Pressure<850)
  - BUT: Analysis requires actual values at those points
- **I/O-bound for value retrieval**
- **Rely on heavy-weight index for fast query process**
  - FastBit: 30%-200% of original data
  - B+ Trees: up to 300% of original data
- **Require access to the original data:**
  - Data + Index

ISABELA-QA: Goals

- Minimize the overhead storage costs,
- Minimize the need for metadata storage in memory,
- Minimize the number of seeks/reads for any particular query,
- Organize the data to allow decomposition of a query into parallel workloads.

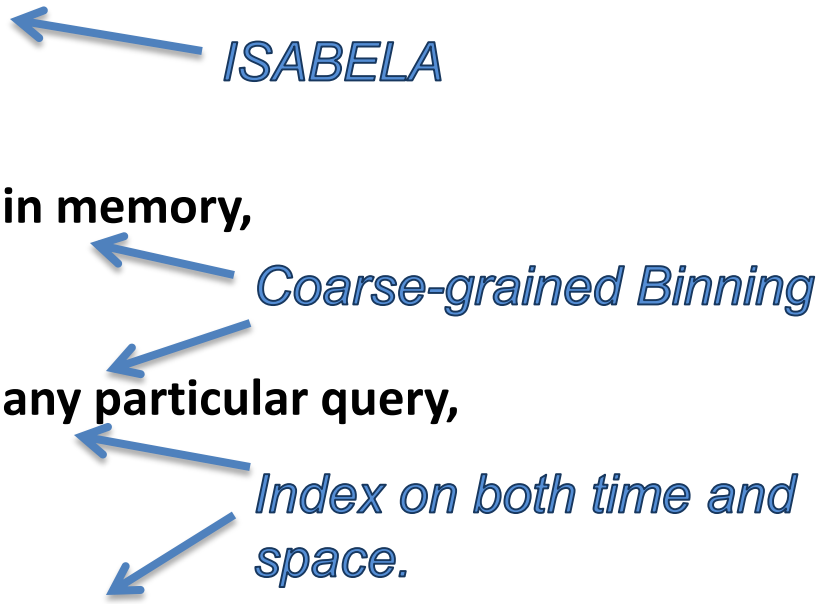

ISABELA-QA: Architecture

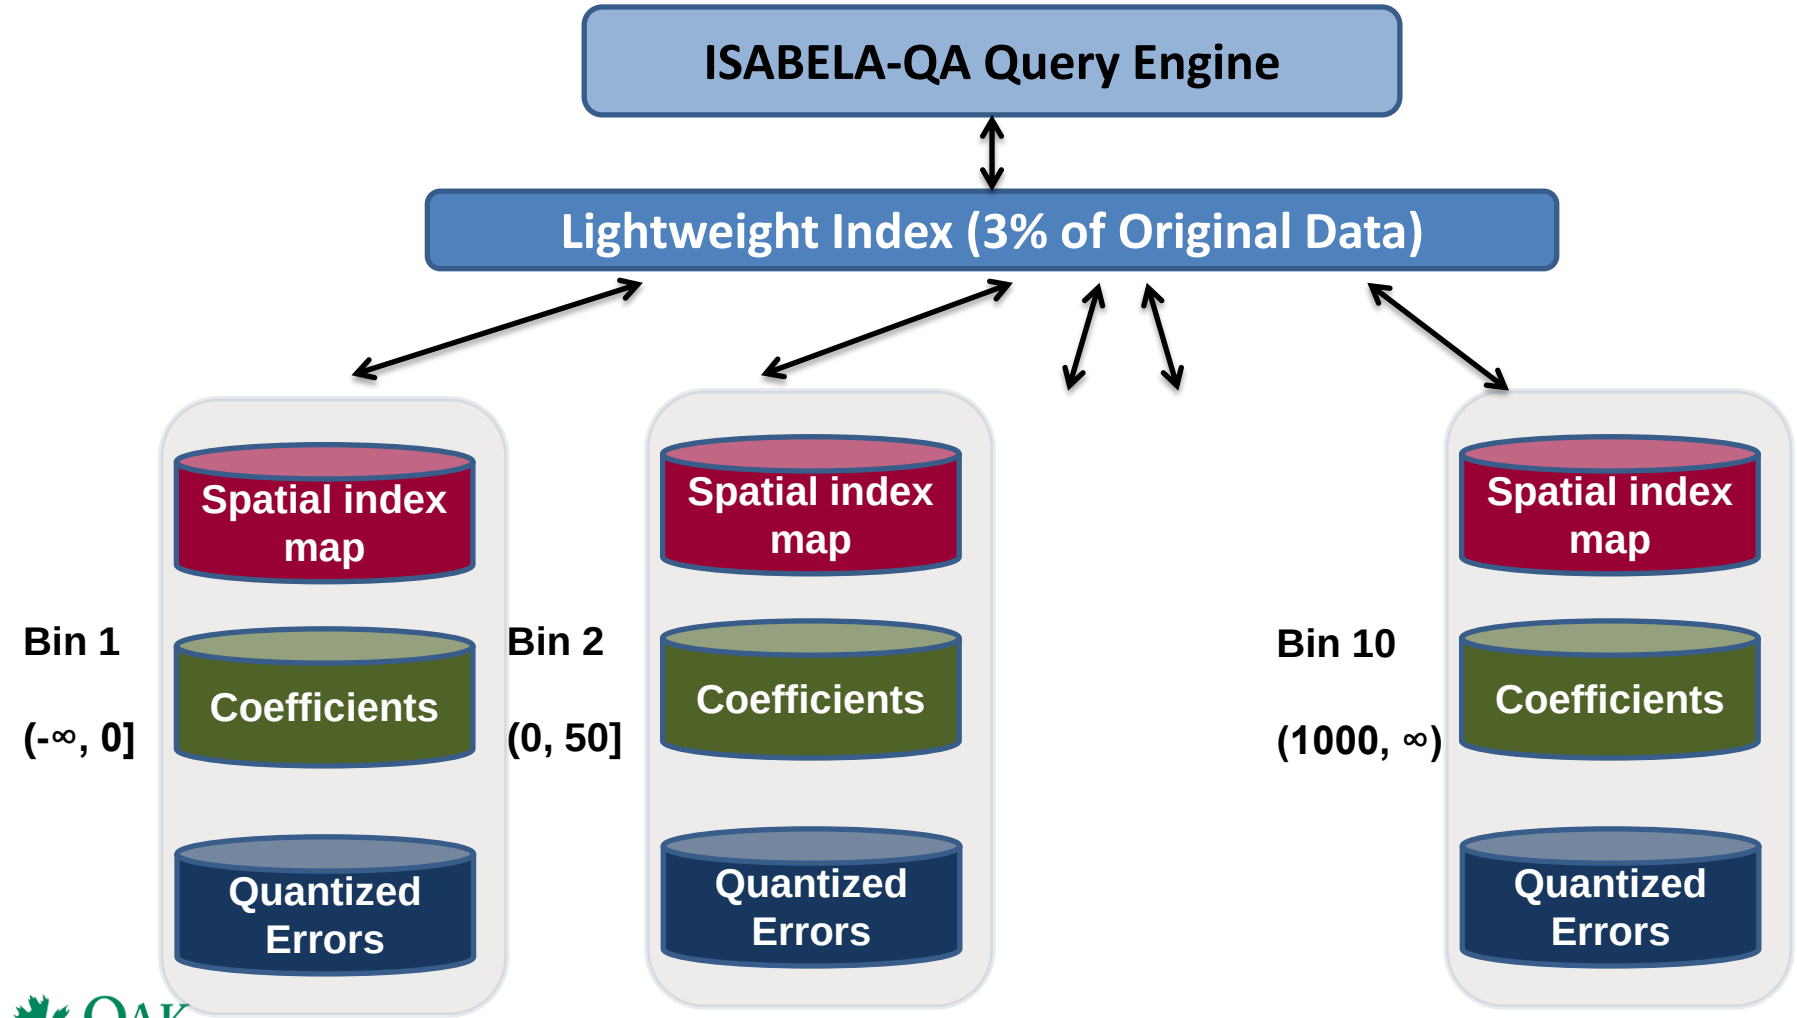

ISABELA-compressed Data (15%-25% of Original Data)

Query Processing Cost with ISABELA-QA

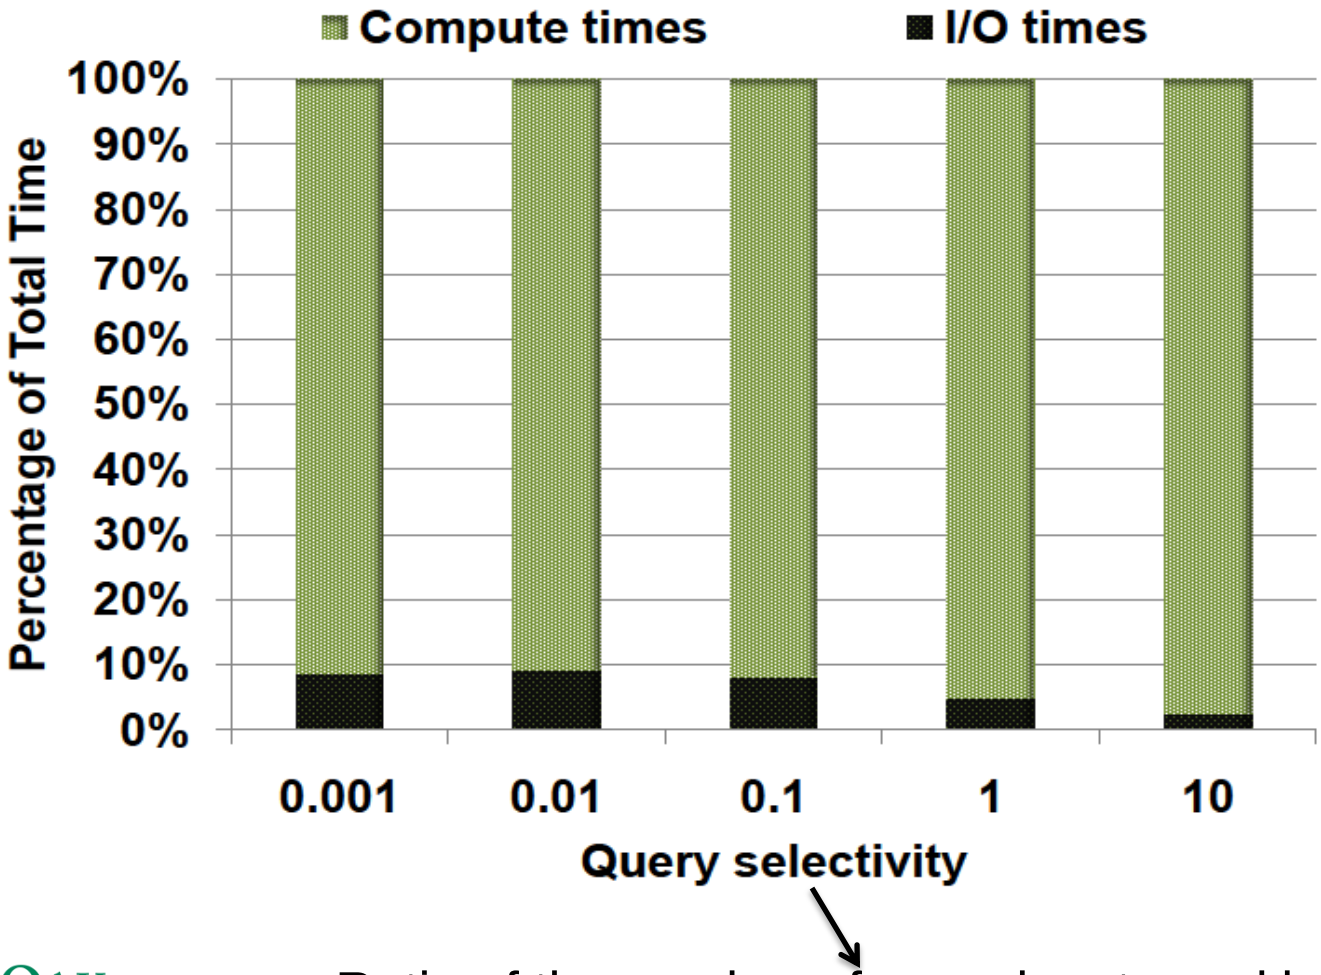

Transition from *I/O*-bound to *GPU/CPU*-bound query processing

Ratio of the number of records returned by the query to the total number of records.

ISABELA-QA: End-to-End Performance

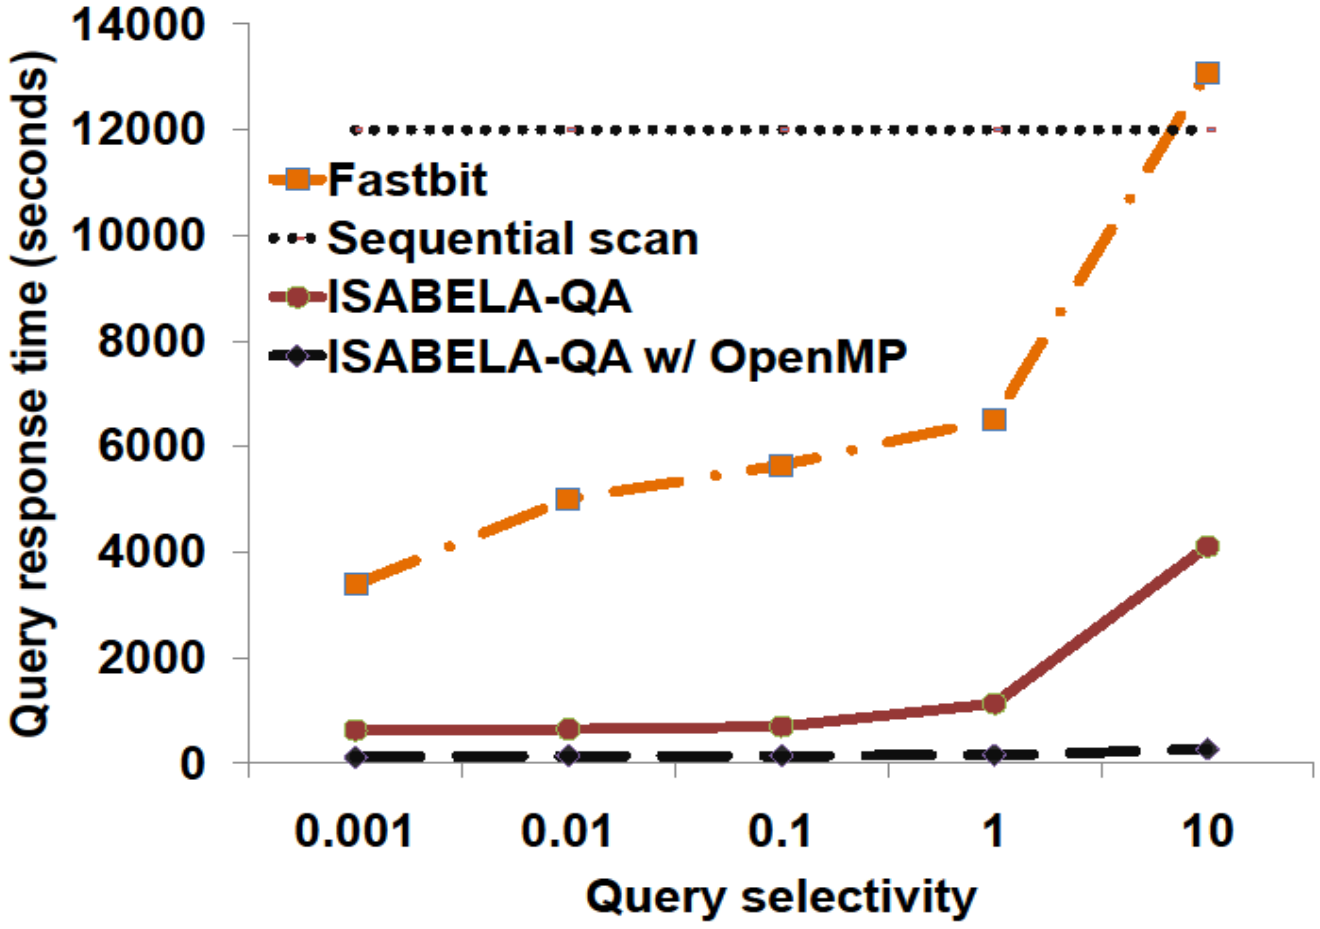

Value-centric queries

- Order of magnitude faster.
- Excellent performance, even for increasing query retrieval.

ISABELA-QA: End-to-End Performance

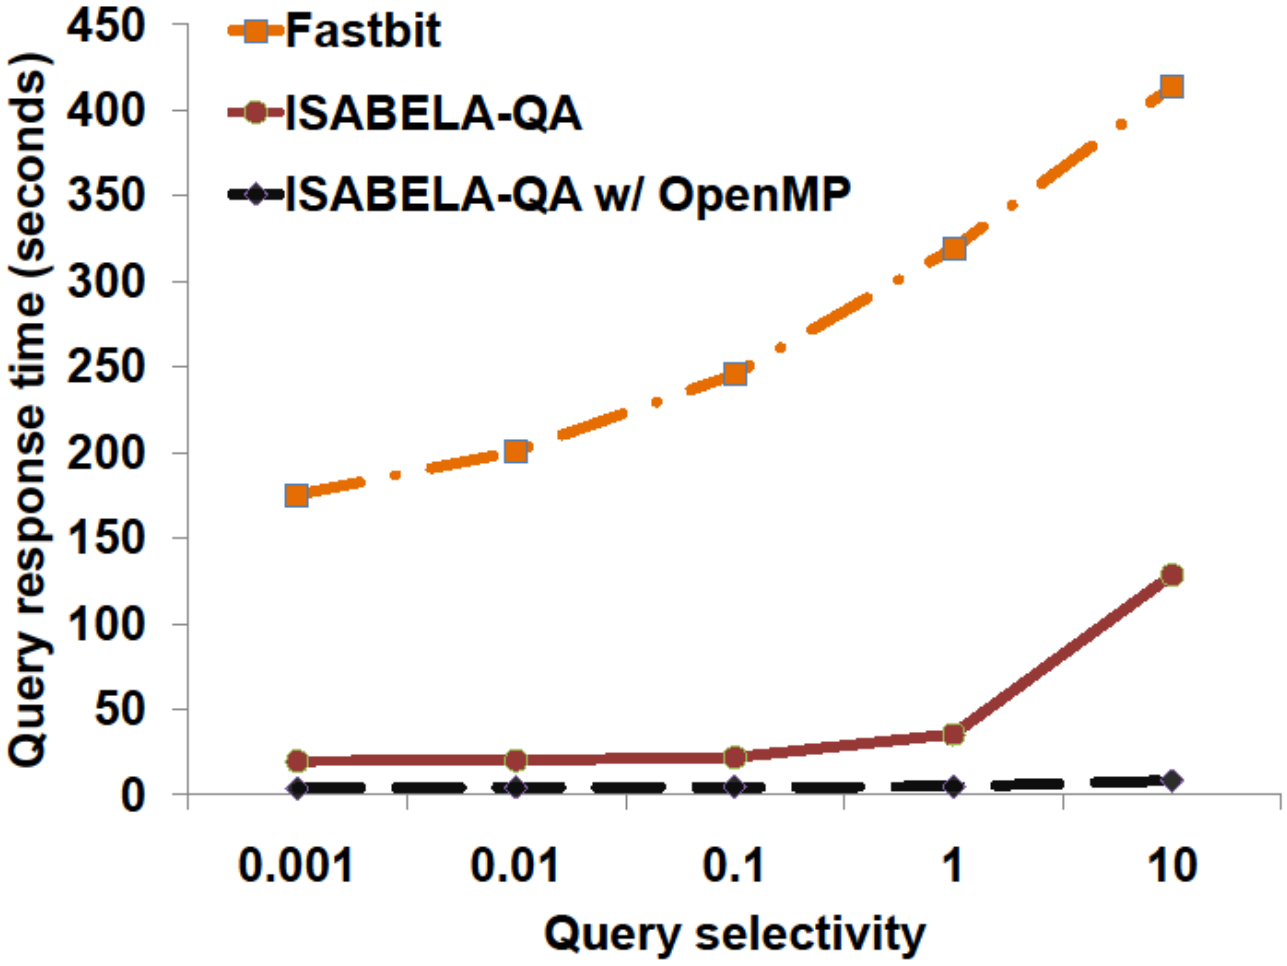

Region-centric queries

ISABELA-QA vs. FastBit Features

Traditional Methods (FastBit)	ISABELA-QA
Optimized for region retrieval	Optimized for both region and especially value retrieval (2-5 orders of magnitude faster)
I/O bound	Trades I/O for CPU/GPU computation
Heavy-weight Index (30%-200% of original data)	Light-weight Index (3% of original data)
Requires original data	Operates over ISABELA-compressed data (15%-20% of original data)
Data + Index = ~270% of original data	Data + Index = 23% of original data
Serial	Embarrassingly Parallel

Speedup with OpenMP

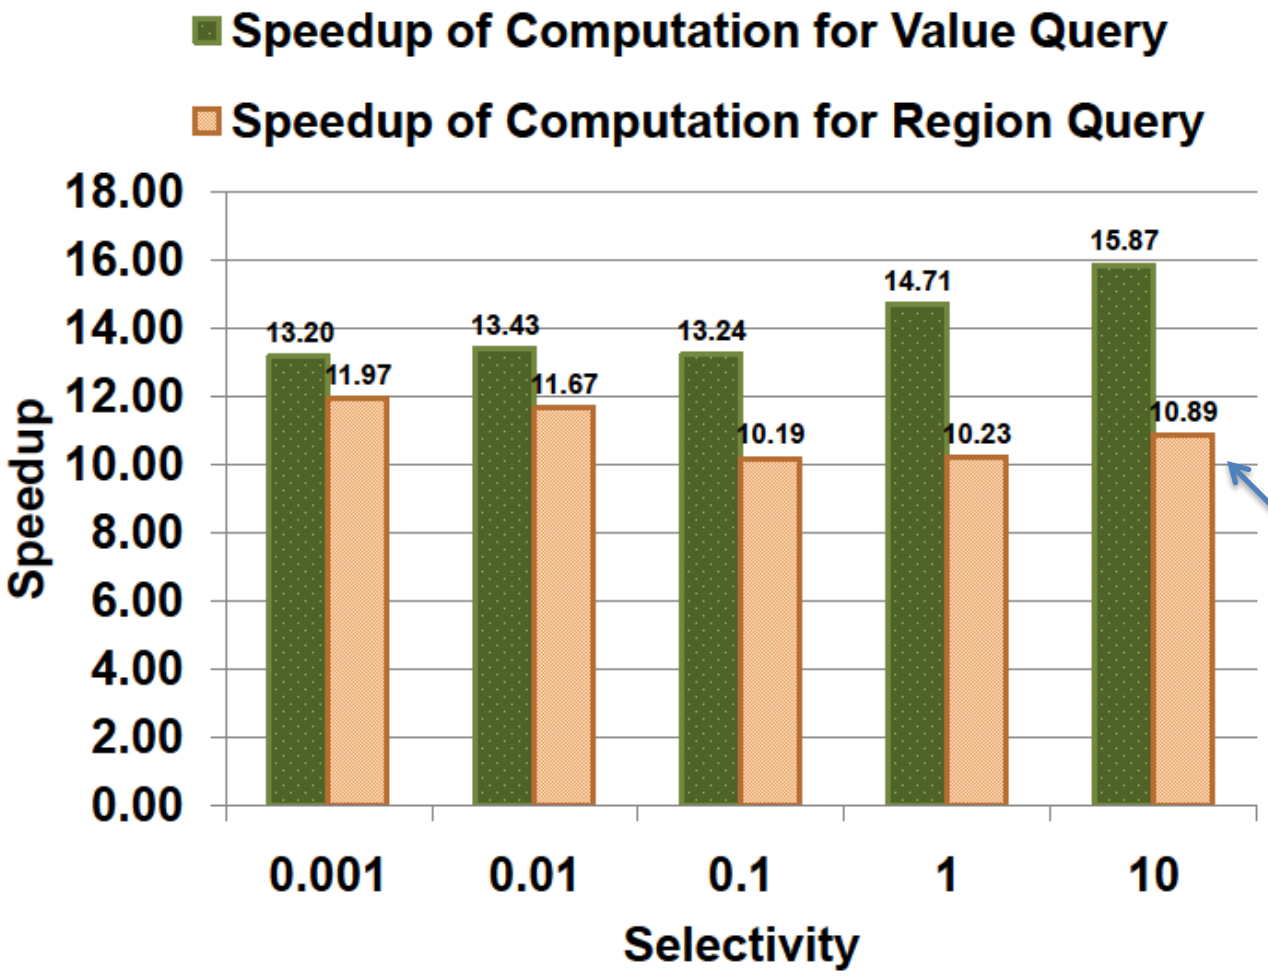

- OpenMP parallelize for timestep.
- 16 OMP Threads

No load balancing was used.

Speedup with GPU

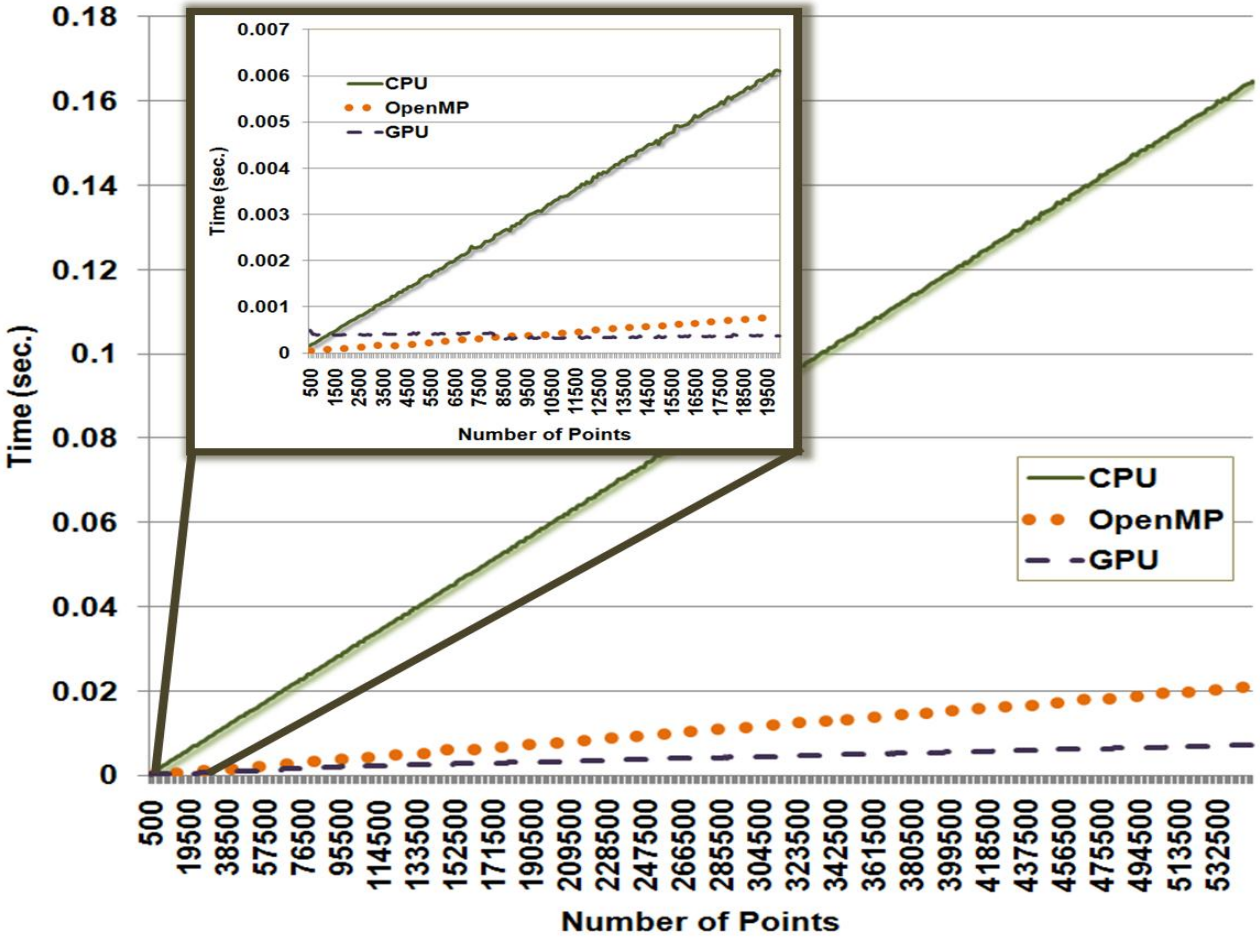

- Embarrassingly parallel. Bspline approximation ideally suited for GPU.