Towards Improving Reverse Time Migration Performance by High-speed Lossy Compression

Yafan Huang (University of Iowa)

05/04/2023

Yafan Huang, Kai Zhao, Sheng Di, Guanpeng Li, Maxim Dmitriev,
Thierry-Laurent D. Tonellot, Franck Cappello

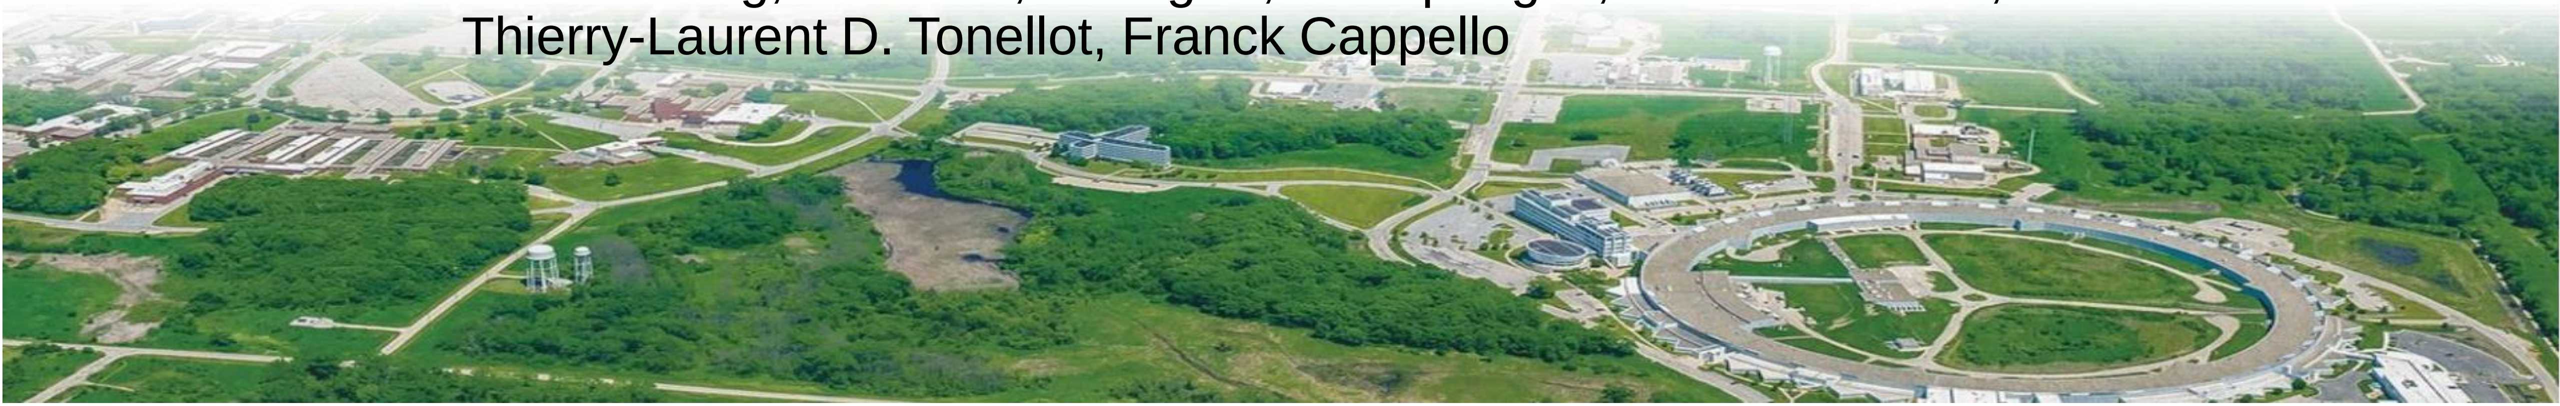

Outline

- Background
- Motivation and Goals
- Related Works
- Methodology
- Evaluation
- Summary

CCGRID 2023
CLUSTER, CLOUD AND INTERNET COMPUTING

Background

- **Seismic imaging** is a procedure for estimating the seismic characteristics of the earth's sub-surface by measuring and computing the acoustic wavefields.
  - Putting vibroseis trucks or air guns and in the ground.
  - Generating acoustic wave data by the sensors.
  - Collecting acoustic wave data and generating earth's subsurface image.

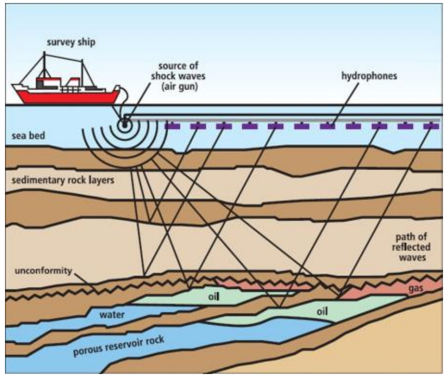

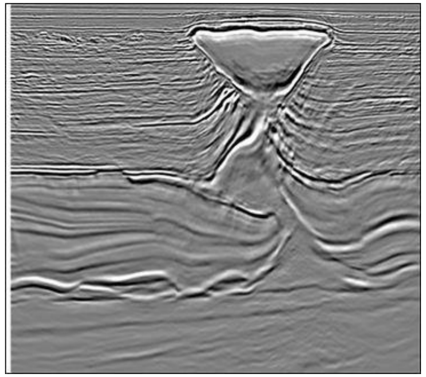

- **Reverse time migration (RTM)** is a cutting-edge seismic imaging method.
  - A full two-way wave equation.
  - Simulating the forward propagation of seismic waves in the subsurface using the wave equation
  - Using the resulting wavefields from backward propagation to create the subsurface image.
  - Generating high resolution image for complex terrain.
  - RTM is widely-used in various domains, such as hydrocarbon exploration, geothermal exploration, and earthquake imaging.

	WEM	Kirchhoff	RTM
Complex velocity focusing?	Yes	No	Yes
Can image steep (>70°) dips?	No	Yes	Yes
Accurate amplitude "out of the box"?	Yes	No	Yes

Limitation of RTM: Very expensive.

U.S. DEPARTMENT OF ENERGY | Office of Science | Argonne NATIONAL LABORATORY | THE UNIVERSITY OF IOWA | UAB THE UNIVERSITY OF ALABAMA AT BIRMINGHAM | أرامكو السعودية saudi aramco

Motivation and Goals

- **Motivation:** Why RTM is so expensive in practice?
  - RTM Workflow:
    1. Data preprocessing (input and configuration)
    2. Forward propagation -> generating snapshots
    3. Backward propagation -> retrieving snapshots
    4. Data postprocessing (generating stacking image)
  - To store the extremely large snapshot data:
    - Peripheral devices: significantly degrading performance.
    - Other nodes: I/O bandwidth, communication overhead.
    - Lossless compression: Limited compression ratio (2:1).

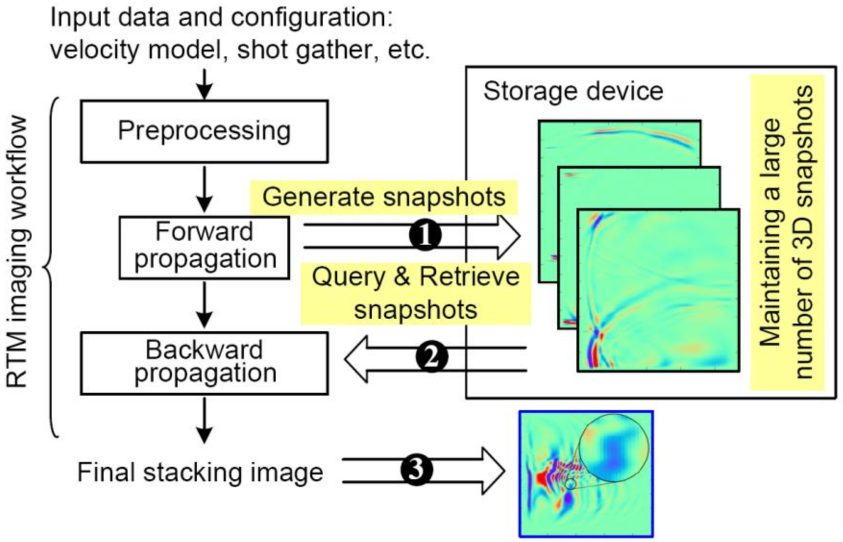

- **Goal:** Error-bounded Lossy Compression with RTM.
  - Drastically reducing the data size while maintain data quality, making it possible to store snapshot data to memory.
  - Improving RTM overall performance.

One snapshot data can be **2800TB** when:
- 10 x 10 km aperture
- 8 km maximum depth

An RTM overall execution have **6k** snapshots:
- Maximum frequency is 80 Hz
- Migration time is 6 seconds

Requirements: **High Throughput** | **High Reconstructed Data Quality** | **Good Compression Ratio for RTM**

U.S. DEPARTMENT OF ENERGY | Office of Science | Argonne NATIONAL LABORATORY | THE UNIVERSITY OF IOWA | UAB THE UNIVERSITY OF ALABAMA AT BIRMINGHAM | أرامكو السعودية saudi aramco

CCGRID 2023
CLUSTER, CLOUD AND INTERNET COMPUTING

Related Works

Resolving Limited Memory in RTM

- **FPGA-based solution** [FPGA'11].
  - Pro: Removing memory constraints.
  - Con: Computation perf. is compromised.
- **Domain-specific data partition** [IPDPS'12].
  - Pros: Paralleling RTM on Blue Gene.
  - Cons: Extra communication overhead.
- **Distributed memory system** [Cluster'19].
  - Pros: Utilizing multiple GPUs.
  - Cons: Data size remains the same.

Error-bounded Lossy Compression

- **ZFP** [TVCG'14]
  - Transform-based lossy compressor.
  - Orthogonal transformation + embedded encoding.
- **SZ** [IPDPS'16, '17]
  - Prediction-based lossy compressor.
  - Prediction + quantization + entropy encoding + lossless compression.

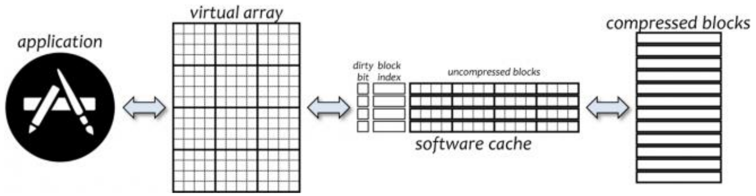

ZFP Workflow

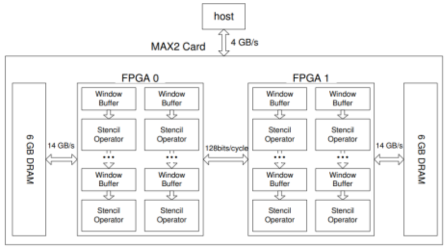

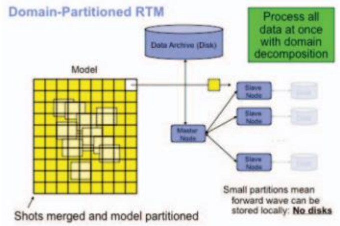

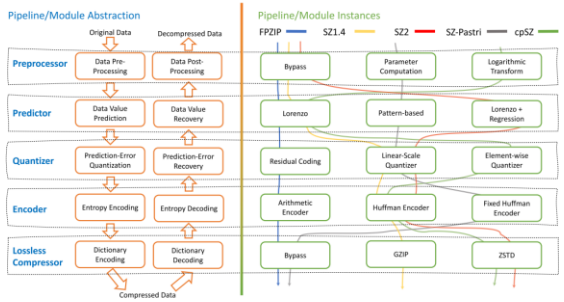

SZ3 Workflow

U.S. DEPARTMENT OF ENERGY | Office of Science
Argonne NATIONAL LABORATORY
THE UNIVERSITY OF IOWA
UAB THE UNIVERSITY OF ALABAMA AT BIRMINGHAM
أرامكو السعودية saudi aramco

Methodology: Outline

- **HyZ** is an error bounded lossy compressor that is specifically designed for RTM.
  - HyZ can satisfy the three key requirements for RTM execution.

High Reconstructed Data Quality | **High Throughput** | **Good Compression Ratio for RTM (5+)**

- We introduce HyZ compressor as follows:
  1. Limitations of SZx, which is an ultra-fast lossy compressor.
  2. BR: Block-wise regression-based lossy compression.
  3. HyZ: Combining BR and SZx for RTM execution.

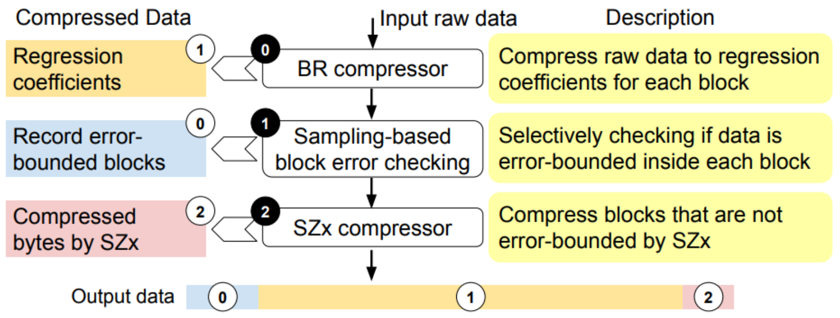

HyZ Workflow

Methodology: Limitations of SZx

- **SZx** [HPDC'18] is one ultra-fast compressor, which may be suitable for RTM execution.
  - SZx Workflow:
    1. Computing the middle value for each block.
    2. **Checking if one block is constant or not.**
    3. Compressing all non-constant blocks.

For constant blocks, only store its middle value.

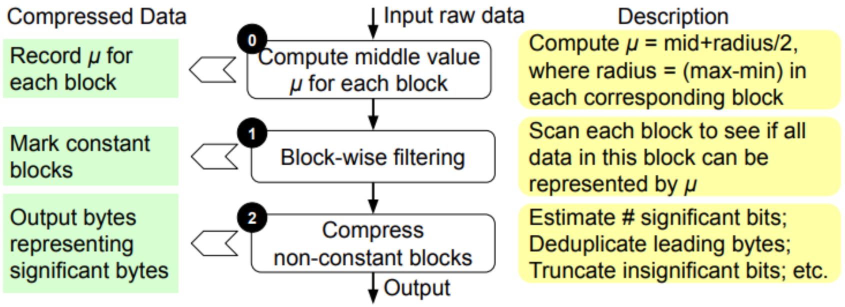

SZx Workflow

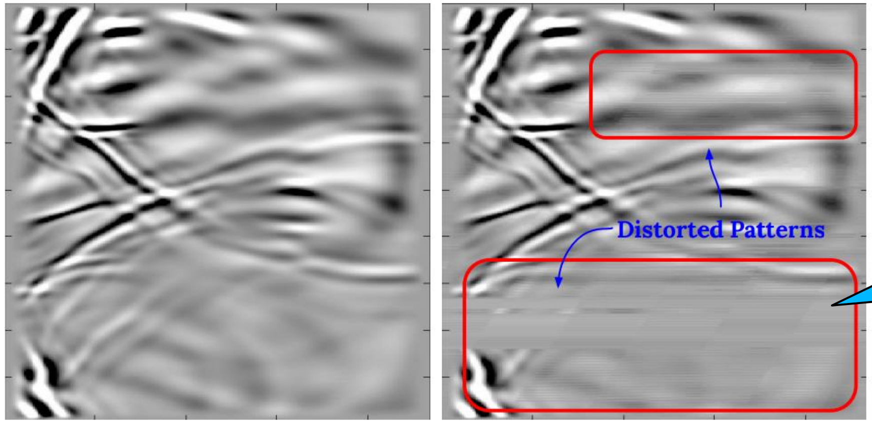

(a) Original Data (b) SZx Reconstructed Data

Visible artifacts! All data in constant blocks is flushed into the same value.

Conclusion: Although SZx is fast, it is not suitable for RTM because of poor reconstructed data quality.

Methodology: BR Compressor

- **BR:** A regression-based ultra-fast lossy compressor.
  - Key insight:
    1. The **linearity** of small data blocks in an RTM snapshot data.
    2. **No dependency** across different blocks, suitable for parallelism
  - BR workflow
    1. Dividing whole dataset into 3D blocks (e.g. 3x3x3).
    2. Predicting each block with a linear hyperplane.

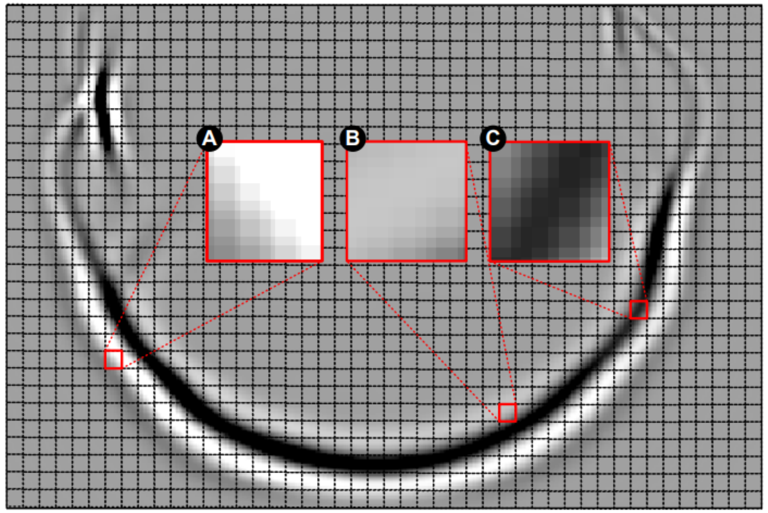

$$f(x,y,z) = \beta_1 x + \beta_2 y + \beta_3 z + \beta_0$$

$$\begin{cases} \beta_1 &= \frac{6}{n_1 n_2 n_3 (n_1+1)}\left(\frac{2V_x}{n_1-1} - V_0\right) \\ \beta_2 &= \frac{6}{n_1 n_2 n_3 (n_2+1)}\left(\frac{2V_y}{n_2-1} - V_0\right) \\ \beta_3 &= \frac{6}{n_1 n_2 n_3 (n_3+1)}\left(\frac{2V_z}{n_3-1} - V_0\right) \\ \beta_0 &= \frac{V_0}{n_1 n_2 n_3} - \left(\frac{n_1-1}{2}\beta_1 + \frac{n_2-1}{2}\beta_2 + \frac{n_3-1}{2}\beta_3\right) \end{cases}$$

where

$$V_0 = \sum_{i=0}^{n_1-1}\sum_{j=0}^{n_2-1}\sum_{k=0}^{n_3-1} d_{ijk}, \quad V_x = \sum_{i=0}^{n_1-1}\sum_{j=0}^{n_2-1}\sum_{k=0}^{n_3-1} i * d_{ijk},$$

$$V_y = \sum_{i=0}^{n_1-1}\sum_{j=0}^{n_2-1}\sum_{k=0}^{n_3-1} j * d_{ijk}, \quad V_z = \sum_{i=0}^{n_1-1}\sum_{j=0}^{n_2-1}\sum_{k=0}^{n_3-1} k * d_{ijk}.$$

Only preserving 4 coefficients for each block data.

Conclusion: Fast enough, good data quality. But it is not error bounded, not enough for RTM post-hoc analysis.

Methodology: HyZ Compressor

- **HyZ:** Combining two ultra-fast lossy compressor BR and SZx.
  - Idea: Designing an ultra-fast compressor with **benefits from both BR and SZx**.
  - HyZ workflow
    1. Compressing raw data by BR compressor.
    2. Checking if a BR-generated block is error-bounded or not.
    3. Compressing non-error-bounded blocks in BR.
  - Analysis for HyZ, why it is suitable for RTM?

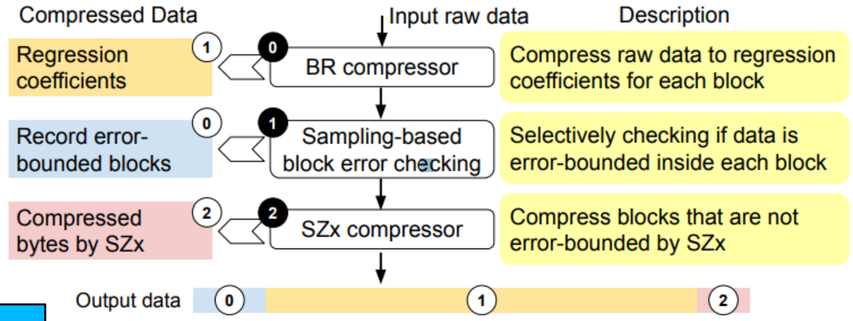

HyZ Workflow

Analysis	Benefit
HyZ is fast, because SZx and BR are ultra-fast designs.	**High Throughput**
HyZ has good compression ratio, since BR only preserve 4 coefficients for each block.	**Good Compression Ratio for RTM (5+)**
HyZ has good reconstructed data quality: introducing error-bounds into BR, which is already a good compressor for RTM execution.	**High Reconstructed Data Quality**

Conclusion: Naturally suitable for RTM execution!

CCGRID 2023
CLUSTER, CLOUD AND INTERNET COMPUTING

Evaluation: Setups

- **Evaluation platform**
  - Bebop supercomputer in Argonne National Laboratory.
    (two Intel Xeon E5-2695 V4 CPUs and 128GB RAM for each node)
- **RTM code and dataset**
  - Industrial parallel RTM code designed by Saudi Aramco.
  - Overthrust belt seismic data (dim: 449x449x235, 3600 snapshots in total).
- **Baseline lossy compressor** (36 threads for each compressor)
  - Transform-based: ZFP [TVCG'14]
  - Prediction-based: SZ [IPDPS'16, '17], SZ-Interp [ICDE'21], SZ-Interp-fast [TBD'23], SZx [HPDC'22]
- **Evaluation metrics**
  - Compression ratio.
  - Compressor speed: compression/decompression throughput (GB/s).
  - Data quality: PSNR, SSIM, visualization data quality.
  - RTM overall performance.

Argonne NATIONAL LABORATORY

أرامكو السعودية
saudi aramco

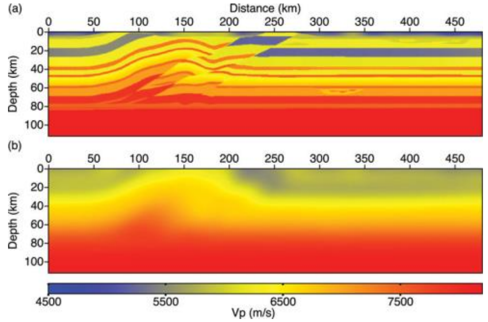

A slice of overthrust data

U.S. DEPARTMENT OF ENERGY | Office of Science | Argonne NATIONAL LABORATORY | THE UNIVERSITY OF IOWA | UAB THE UNIVERSITY OF ALABAMA AT BIRMINGHAM | أرامكو السعودية saudi aramco

Evaluation: Compression Ratio

Compressor	Relative Error Bound 1E-3	Relative Error Bound 4E-4	Absolute Error Bound 1E-5	Absolute Error Bound 4E-6	**Average**
SZx	7.58	6.22	9.57	7.72	7.77
SZ-Interp	213.53	116.60	177.92	113.20	155.31
SZ-Interp-fast	17.05	9.09	16.26	11.06	13.37
SZ	26.94	25.51	26.83	24.37	25.91
SZ(Serial)	70.36	50.68	81.47	54.94	64.36
ZFP	26.92	20.65	38.74	32.04	29.59

Compressor	Block Size $3 \times 3 \times 3$	$4 \times 4 \times 4$	$5 \times 5 \times 5$	$6 \times 6 \times 6$	**Average**
BR	6.66	15.72	31.11	52.64	26.53

Compression ratio for other compressors.

HyZ	Relative Error Bound 1E-3	Relative Error Bound 4E-4	Absolute Error Bound 1E-5	Absolute Error Bound 4E-6	**Average**
$3 \times 3 \times 3$	5.01	4.61	5.17	4.99	12.31
$4 \times 4 \times 4$	9.14	8.20	10.29	9.69	
$5 \times 5 \times 5$	13.87	12.24	17.10	15.59	
$6 \times 6 \times 6$	18.50	16.11	24.84	21.52	

Compression ratio for HyZ.

Takeaway 1: HyZ and BR both can meet compression ratio requirements (5+) in RTM execution across 3600 time steps, with 12.31 and 26.53 on average.

Evaluation: Compression Speeds

Compressor	Relative Error Bound 1E-3	Relative Error Bound 4E-4	Absolute Error Bound 1E-5	Absolute Error Bound 4E-6	**Average**
SZx	12.48	11.08	16.00	8.78	12.09
SZ-Interp	2.87	2.76	2.76	2.47	2.72
SZ-Interp-fast	5.48	5.36	5.53	5.11	5.37
SZ	3.63	3.57	3.61	3.57	3.60
ZFP	2.82	2.56	4.33	4.04	3.44

Compressor	Block Size $3 \times 3 \times 3$	$4 \times 4 \times 4$	$5 \times 5 \times 5$	$6 \times 6 \times 6$	**Average**
BR	5.85	15.64	24.36	32.46	19.58

Compression speeds
for other compressors

Compressor	Relative Error Bound 1E-3	Relative Error Bound 4E-4	Absolute Error Bound 1E-5	Absolute Error Bound 4E-6	**Average**
SZx	19.65	18.13	22.45	15.44	18.92
SZ-Interp	5.26	4.88	4.77	4.48	4.85
SZ-Interp-fast	9.00	9.10	9.61	8.03	8.94
SZ	5.14	4.75	5.01	4.76	4.92
ZFP	0.60	0.54	0.62	0.58	0.59

Compressor	Block Size $3 \times 3 \times 3$	$4 \times 4 \times 4$	$5 \times 5 \times 5$	$6 \times 6 \times 6$	**Average**
BR	10.67	19.36	24.96	26.59	20.40

Decompression speeds
for other compressors

HyZ	**Compression Speed** Relative Error Bound 1E-3	4E-4	Absolute Error Bound 1E-5	4E-6	**Average**	**Decompression Speed** Relative Error Bound 1E-3	4E-4	Absolute Error Bound 1E-5	4E-6	**Average**
$3 \times 3 \times 3$	4.54	4.28	4.72	4.33	10.69	7.45	7.33	7.65	7.26	12.45
$4 \times 4 \times 4$	8.45	8.37	8.83	8.34		12.71	12.16	12.94	11.83	
$5 \times 5 \times 5$	13.12	11.93	13.42	12.32		14.64	14.43	15.01	13.47	
$6 \times 6 \times 6$	17.64	16.41	17.94	16.40		15.85	15.47	16.01	15.03	

Compression and decompression speeds for HyZ

Takeaway 2: HyZ and BR both have top-class compression and decompression speeds. HyZ can achieve 10.69 GB/s and 12.45 GB/s for compression and decompression.

Evaluation: Data Quality

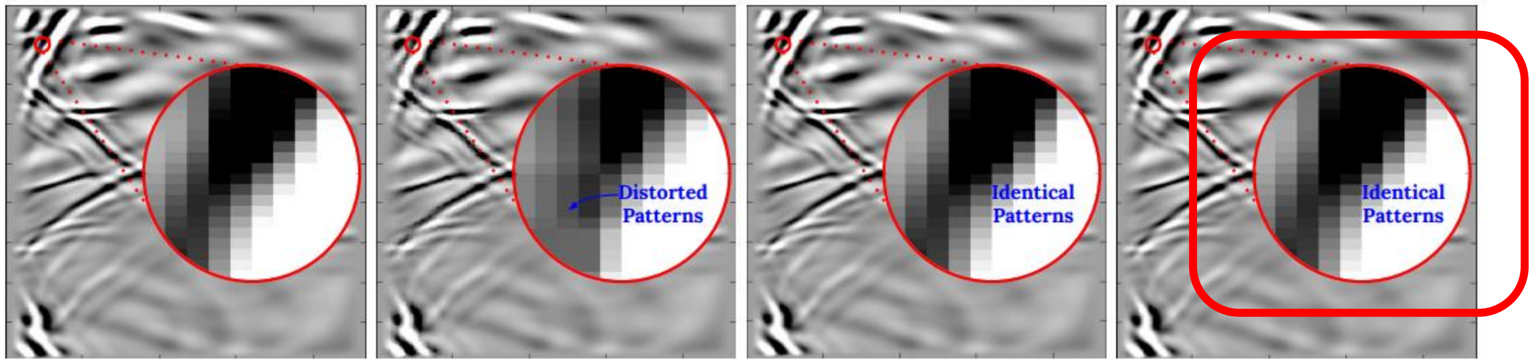

(a) Original Data

(b) BR, Block Size=5 × 5 × 5, (PSNR: 50.86dB, SSIM: 0.9546)

(c) HyZ, Block Size=5 × 5 × 5, REL=1E-3 (PSNR: 87.91dB, SSIM: 0.9980)

(d) HyZ, Block Size=6 × 6 × 6, REL=1E-3 (PSNR: 68.78dB, SSIM: 0.9861)

Single snapshot (snapshot=3000)

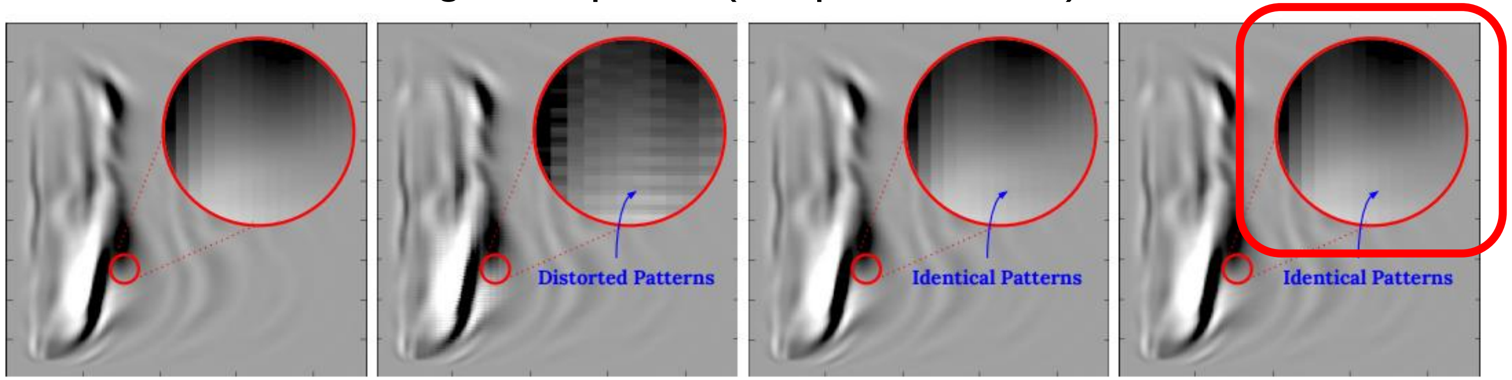

(a) Original Data

(b) BR, Block Size=4 × 4 × 4, (PSNR: 74.50dB, SSIM: 0.9713)

(c) HyZ, Block Size=4 × 4 × 4, (PSNR: 124.13dB, SSIM: 0.9935)

(d) HyZ, Block Size=5 × 5 × 5, (PSNR: 117.35dB, SSIM: 0.9902)

Final stacking image

Takeaway 3: HyZ is superior to BR in preserving the data quality in both single snapshot and final stack imaging, also with higher PSNR and SSIM.

Evaluation: RTM Overall Performance

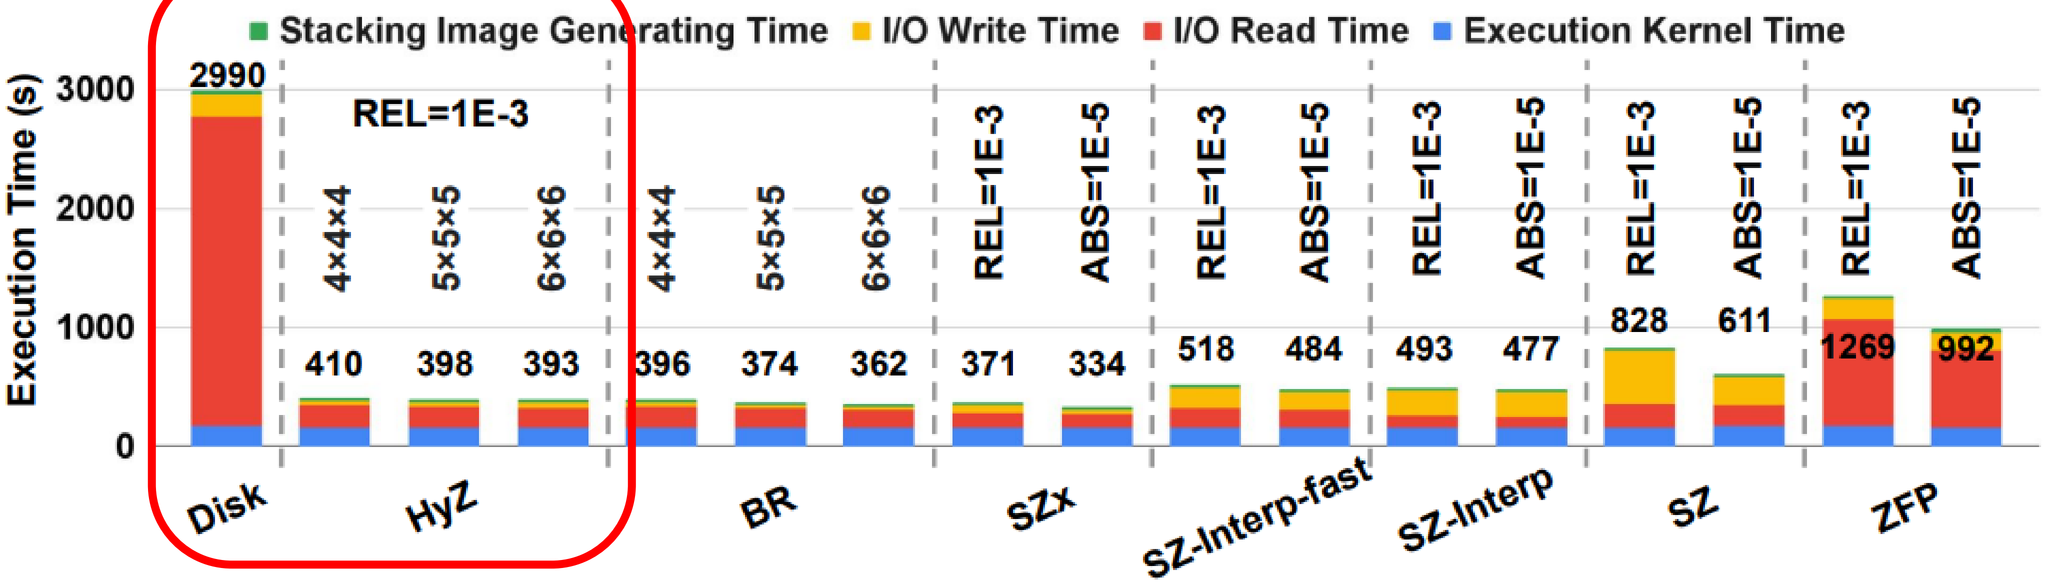

Takeaway 4: HyZ and BR both can lead to a top-tier performances in RTM execution, increasing the performance on average by 6.46x and 6.93x, compared with RTM without lossy compression.

Summary

- **Conclusion**
  - We use lossy compression to resolve the memory capacity issue in RTM execution.
  - We design HyZ, a domain-specific error-bounded compressor, for RTM execution.
  - Evaluation demonstrates HyZ is the best lossy compressor for RTM, with high speeds, high compression, and high reconstructed data quality.
- **Future works**
  - Accelerating RTM in GPU computing environments.
  - Optimizing HyZ performance.

Yafan Huang
University of Iowa
yafan-huang@uiowa.edu
https://hyfshishen.github.io/

This research was supported by the U.S. Department of Energy, Office of Science, Advanced Scientific Computing Research (ASCR), under contract DE-AC02-06CH11357, and supported by the National Science Foundation under Grant OAC-2003709, OAC-2104023, OAC-2211538, and OAC2211539/2247060. We acknowledge the computing resources provided on Bebop (operated by Laboratory Computing Resource Center at Argonne) and on Theta and JLSE (operated by Argonne Leadership Computing Facility). This research was also supported by ARAMCO.