أرامكو السعودية
saudi aramco

Lossless Compression of Internal Files in Parallel Reservoir Simulation

Suha Kayum
Marcin Rogowski
Florian Mannuss

9/26/2019

where energy is opportunity™

Outline

- I/O Challenges in Reservoir Simulation
- Evaluation of Compression Algorithms on Reservoir Simulation Data
- Real-world application
  - Constraints
  - Algorithm
  - Results
- Conclusions

Challenge

Reservoir simulation

1

Reservoir Simulation

- Largest field in the world are represented as 50 million – 1 billion grid block models
- Each runs takes hours on 500-5000 cores
- Calibrating the model requires 100s of runs and sophisticated methods
- "History matched" model is only a beginning

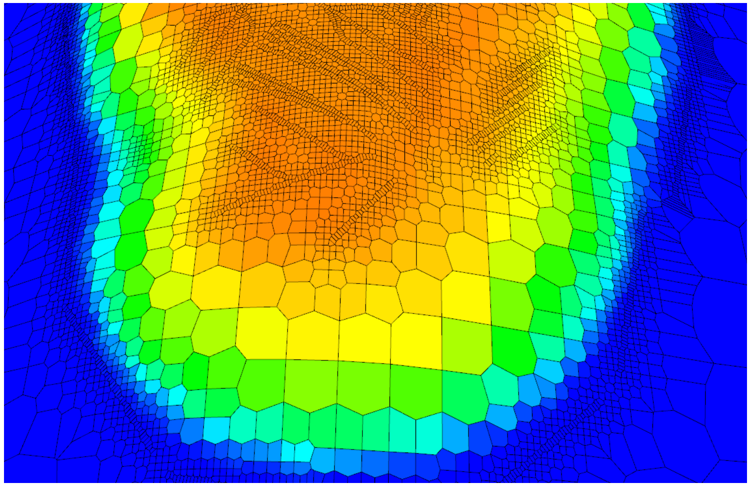

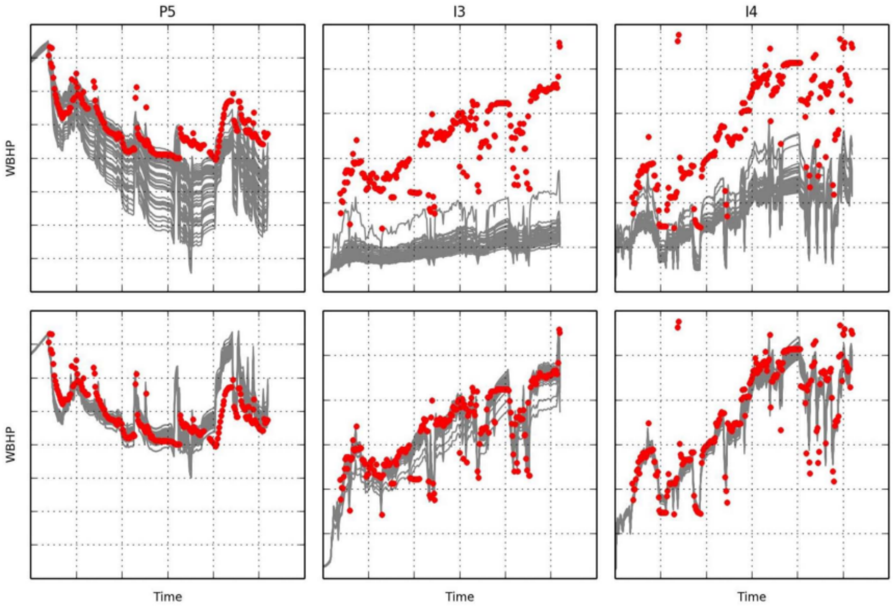

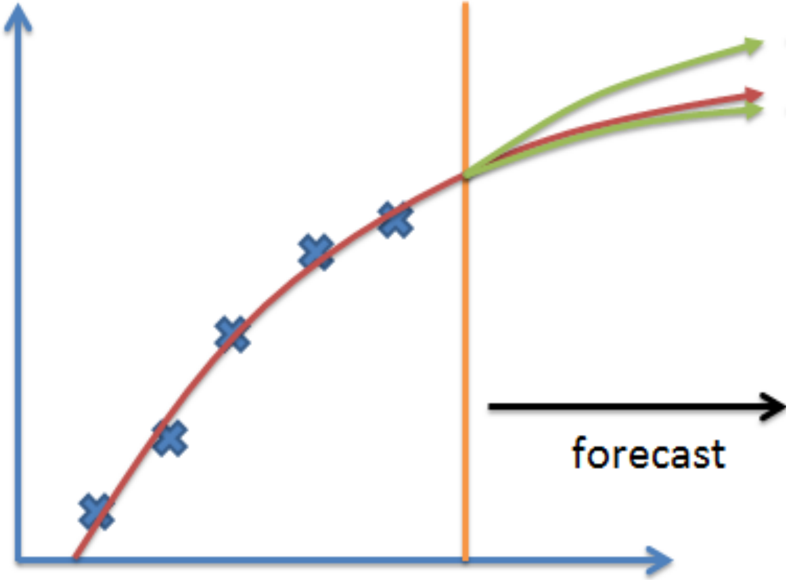

Files in Reservoir Simulation

- Internal Files
- Input / Output Files
  - Interact with pre- & post-processing tools

$$\frac{\partial c_i}{\partial t} + \nabla.\sum_{j=1}^{n_p}\left(\rho_j\omega_{ij}u_j - \phi\rho_j S_j \vec{D}_{ij}.\nabla\omega_{ij}\right) = \phi\sum_{j=1}^{n_p}\left(S_jR_{ij}\right) + (1-\phi)R_{is} + q_i$$

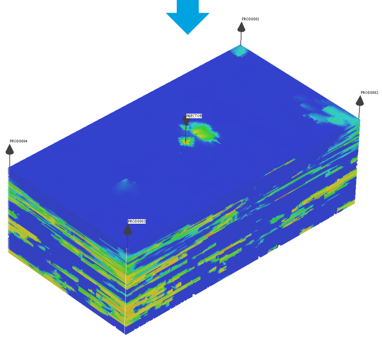

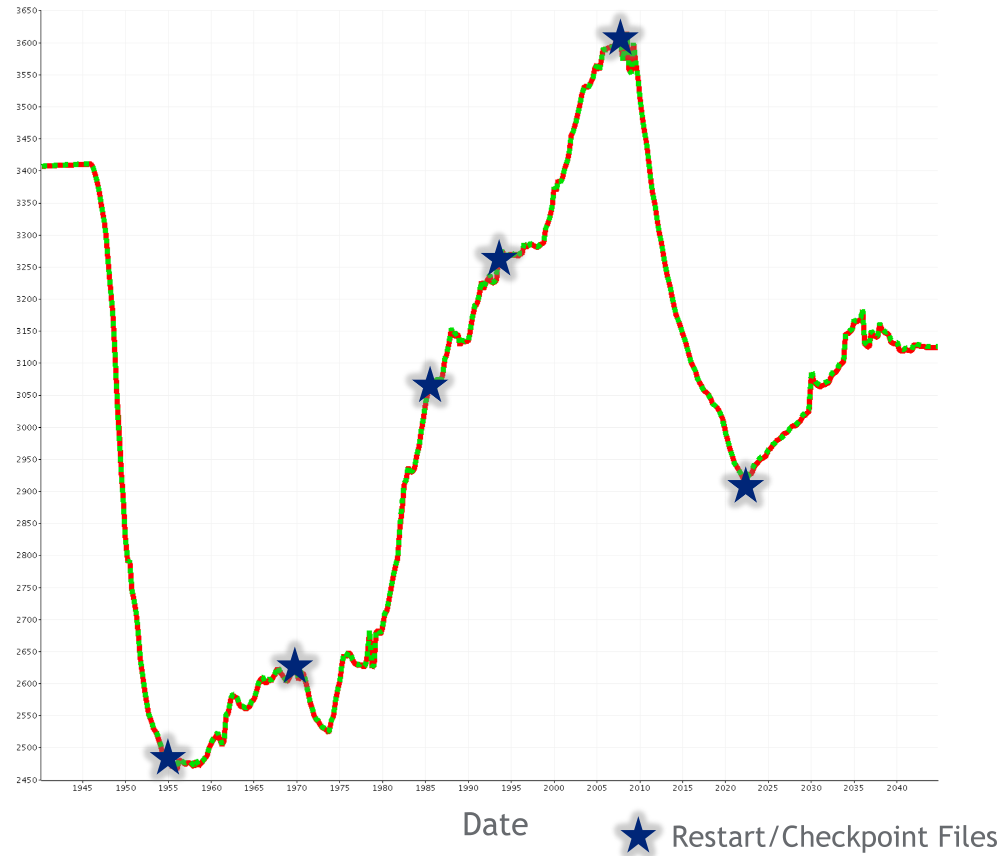

Reservoir Simulation in Saudi Aramco

- 100’000+ simulations annually
- The largest simulation of 10 billion cells
- Currently multiple machines in TOP500
- Petabytes of storage required
- Resources are Finite
- File Compression is one solution

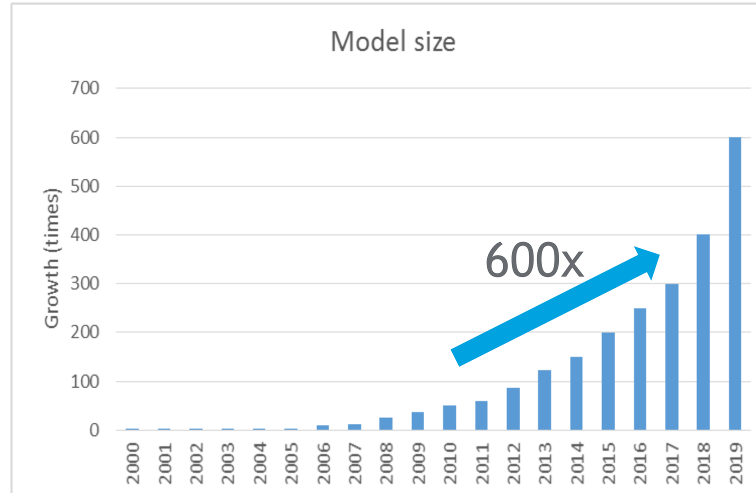

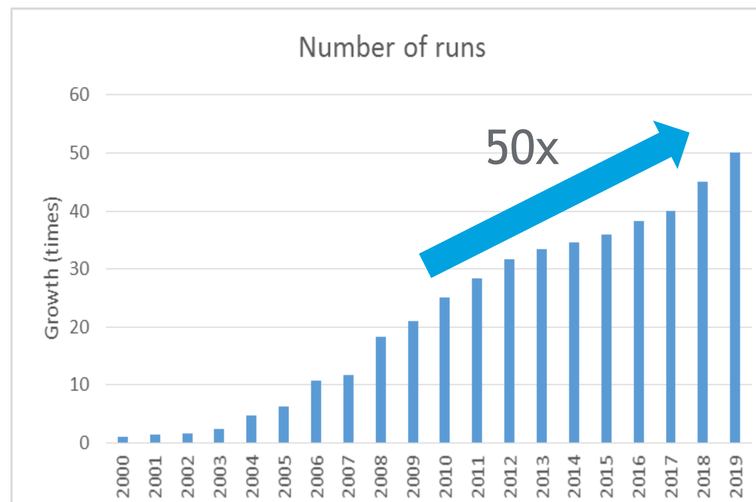

Compression algorithm evaluation

2

Compression ratio

Tested a number of algorithms on a GRID restart file for two models

- Model A – 77.3 million active grid blocks
- Model K – 8.7 million active grid blocks
- 15.6 GB and 7.2 GB respectively

Compression ratio is between

- From 2.27 for snappy (Model A)
- Up to 3.5 for bzip2 -9 (Model K)

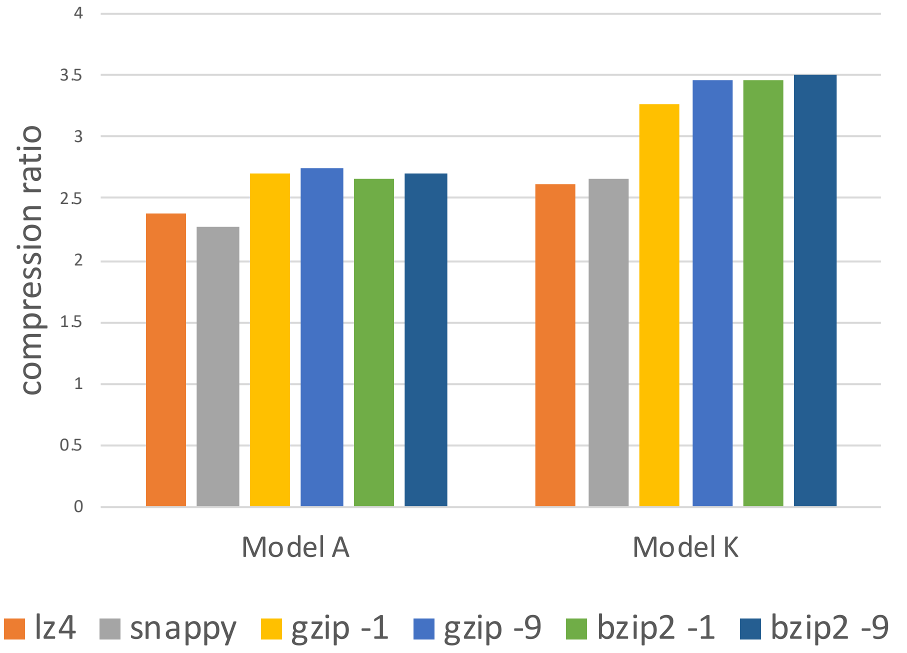

Compression speed

- LZ4 and Snappy significantly outperformed other algorithms when it comes to runtime overhead

- Compression ratio of LZ4 was 14-34% worse than bzip2 variant
  - Overhead was 117-138 times smaller when it comes to compression
    - 34s vs 4680s for 15.6GB
  - Overhead 19-23 times smaller when decompressing

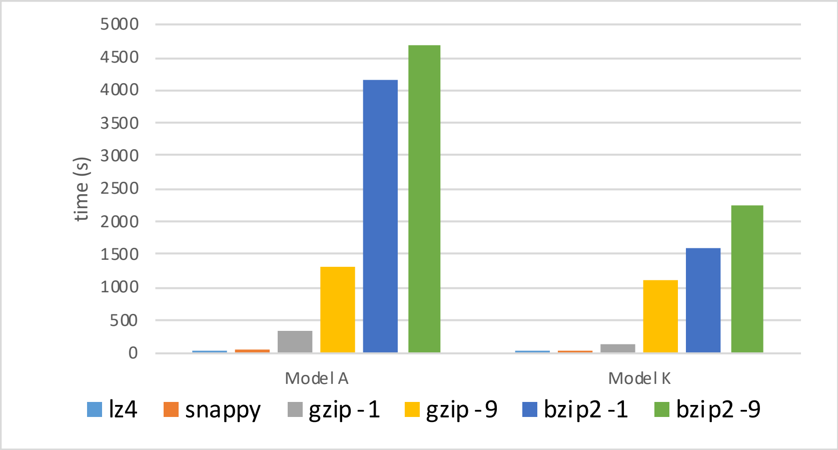

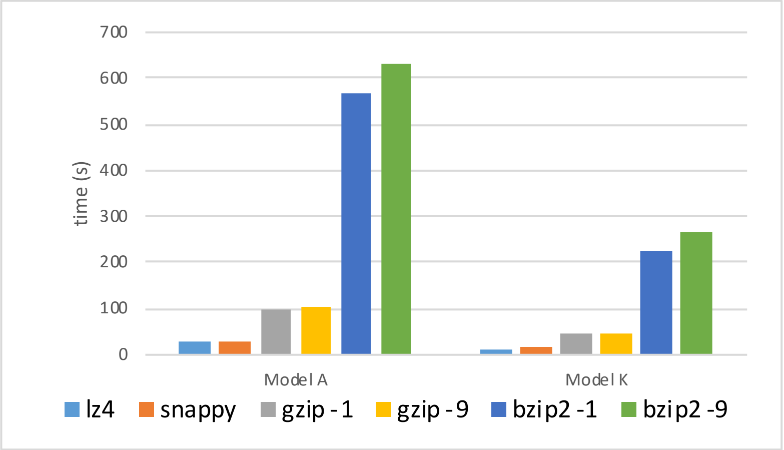

Real-world

3

Constraints

- Fitting within a framework of a large, pre-existing application
- Limiting the number of files
- Parallel implementation
- Compression ratio – runtime overhead balance

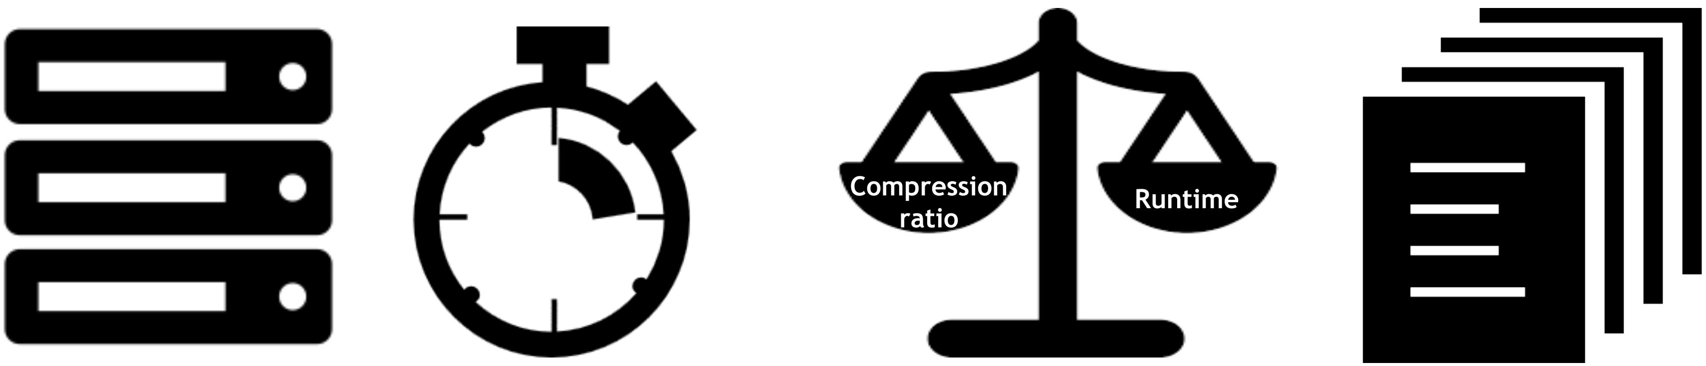

Algorithm

Compression Method 1 Double or integer arrays interpreted directly as character arrays

Compression Method 2 (Default)

- Arrays are reordered where all first bytes of double-precision values are stored in one array, the second bytes in another one and so on.
- Create 8 arrays, each corresponding to the N^{th} byte of every floating-point value
- Reduces the entropy of sub-arrays that might be introduced due to the IEEE floating-point standard
- 8 arrays are compressed individually

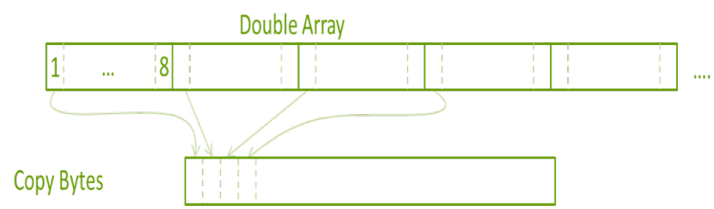

Illustration of array reordering with the aim to reduce entropy.

Indexing

- After compression, writer processes exchange compressed array sizes and write all compressed chunks to file
- An index file is created by writers sending their chunk sizes to the writer with lowest rank

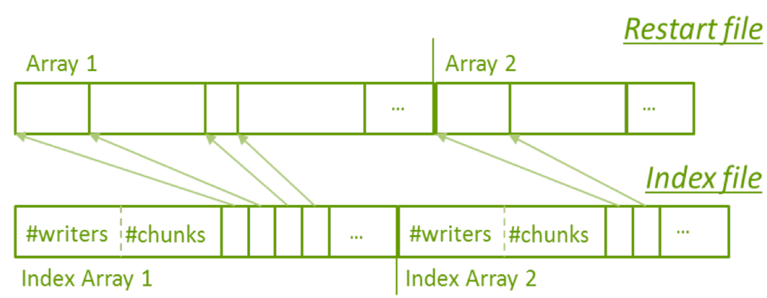

Index structure pointing to the start of the individual compressed chunks in the restart file.

Real-world results

- Tested on 500+ synthetic and real-field models that encompass varying model features and sizes
- Lowest compression efficiency resulted in a ratio of 1.47
- Best-case scenario, the ratio exceeded 25
- The median model compressed with a ratio of 3 with the mean at 3.6.
- Cumulative size of restart GRID files was reduced from
  - 680 GB to 232 GB
  - compression ratio of 2.93

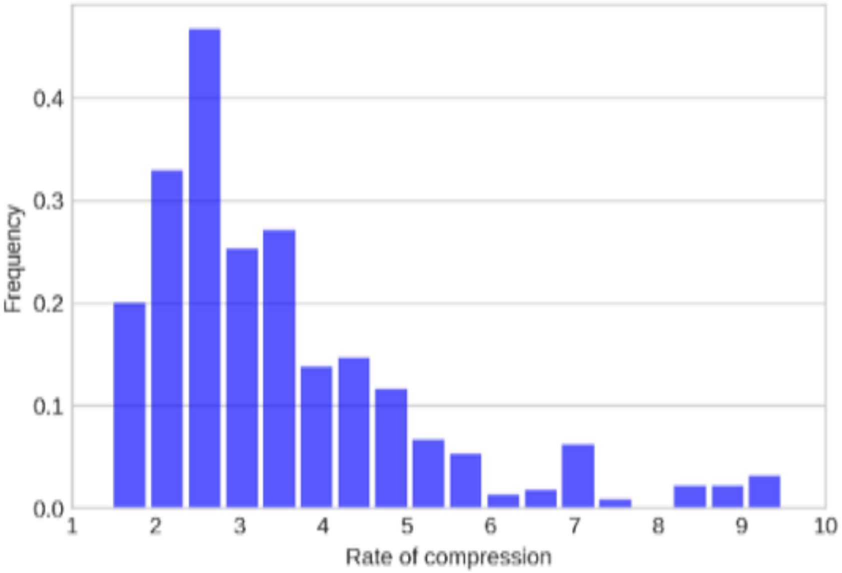

أرامكو السعودية
saudi aramco

Conclusions

- Storage and performance challenges of recurrently writing massive restart files in a parallel reservoir simulator are described and addressed using compression
- Evaluated a number of lossless compression algorithms
- LZ4 and Snappy performed the best in terms of compression speed resulting in less runtime overhead (LZ4 chosen)
- Compression was implemented in parallel
- Reduced storage requirements and increased efficiency
  - In some cases, the overall runtime was reduced, i.e. compressing and writing to disk is faster than writing raw data to disk
- Shifting the load from interconnect and storage to performing additional computation is in line with predicted trends in the high-performance computing where compute capabilities are expected to grow at a higher rate than any other supercomputer component

أرامكو السعودية
saudi aramco