Modeling power usage for the SZ lossy compressor on HPC systems

Grant Wilkins, Jon Calhoun (Advisor)
Holcombe Department of Electrical and Computer Engineering

Introduction

Motivation:
Lossy compressors (e.g. SZ [2] and ZFP [9]) are increasing in ubiquity as they are more compress floating-point data better.

As scientific simulations become more complex and create more data, more energy will be spent in compressing data due to runtime increases [4].

Understanding how variables such as CPU frequency, error bound, and data size affect power will increase the ability to save energy.

Contributions:
- We present a characterization of nodal energy consumption for SZ compression and decompression.
- We provide a predictive model for system power given CPU frequency and, minimum frequency power.
- We produce recommendations for optimizing energy consumption using the gathered results from the energy model.

Experimental Setup

Application	Dimensions	Size of Fields	Fields Considered
CESM-ATM [6, 8]	$26 \times 1800 \times 3600$	673.9MB	CLOUD, CMFDT, ICLDTWP, QC, T, VU
HACC [5, 8]	1×280953867	1046.9MB	v_x, v_y, v_z, x, y, z
NYX [7, 8]	$512 \times 512 \times 512$	563.9MB	b_dens, dm_dens, temp, v_x, v_y, v_z

Table 1. Datasets used in SZ power monitoring

Intel "Skylake" Xeon Silver 4114 CPUs used in every test.

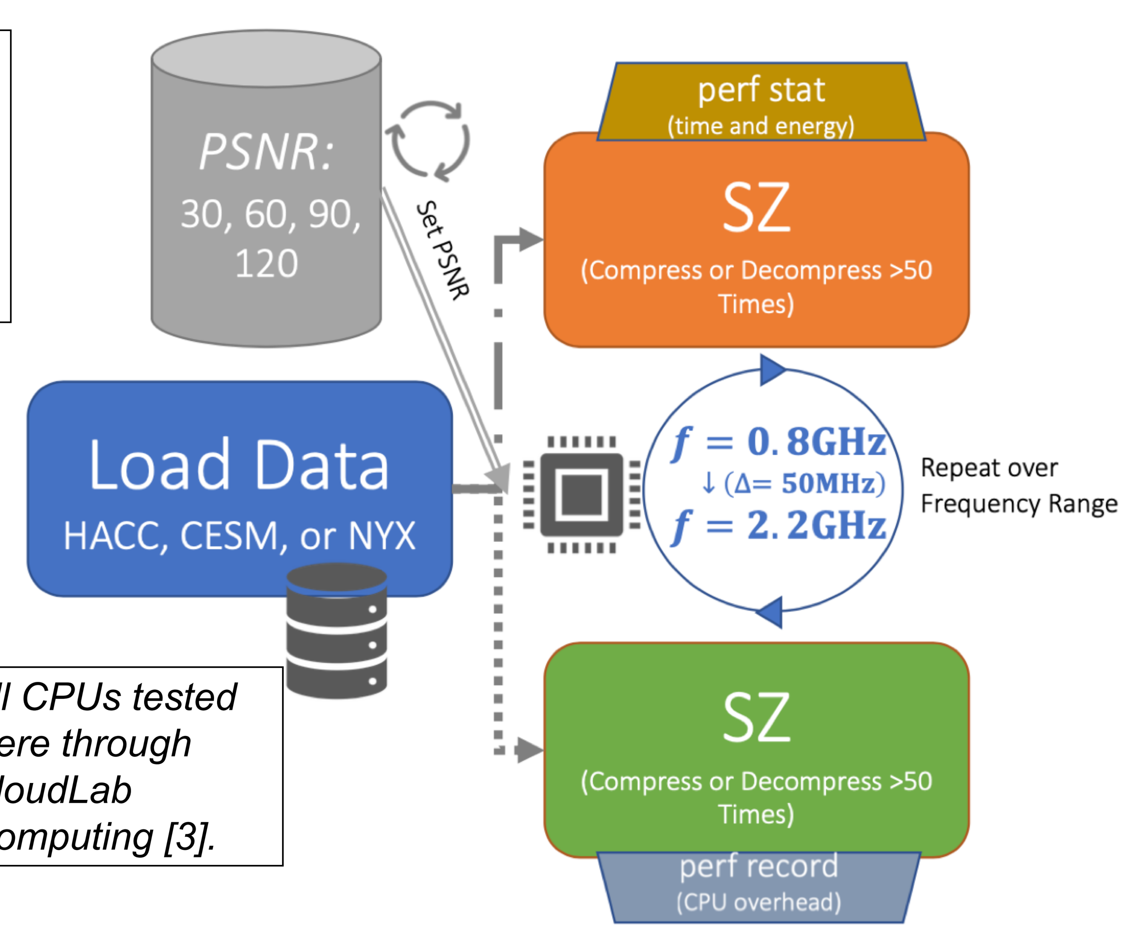

All CPUs tested were through CloudLab Computing [3].

Link to Personal Website

Fig. 1 Experimental process for SZ monitoring

Making a Power Approximation

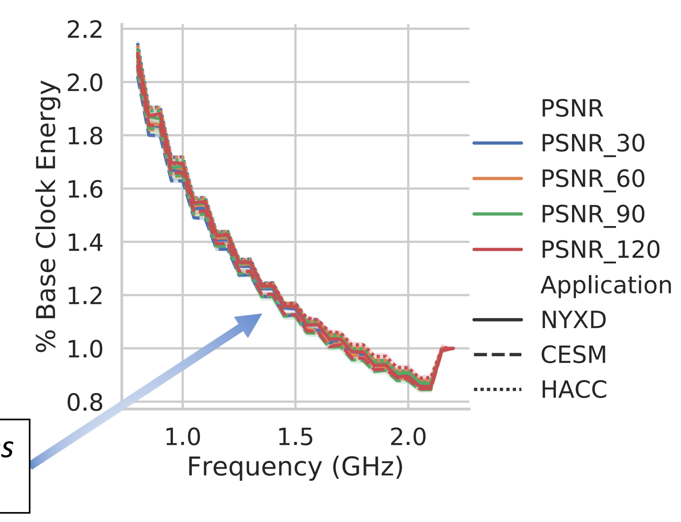

CESM, HACC, NYX curves indistinguishable

Fig. 2 SZ Energy

CESM, HACC, NYX curves overlapping

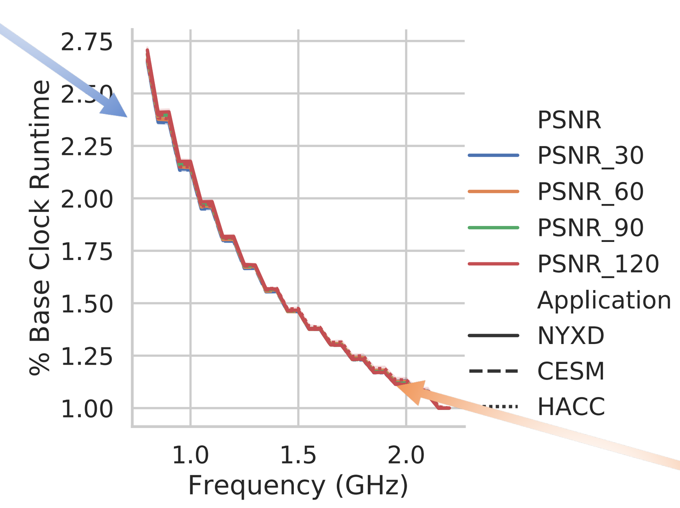

Fig. 3 SZ Runtime

$$P_{avg} = \frac{E_{total}}{t_{SZ}}$$

Eqn. 1 Average power relation

*No significant difference between **PSNR** or **Application** in profile*

Model computed from datasets in Table 1.

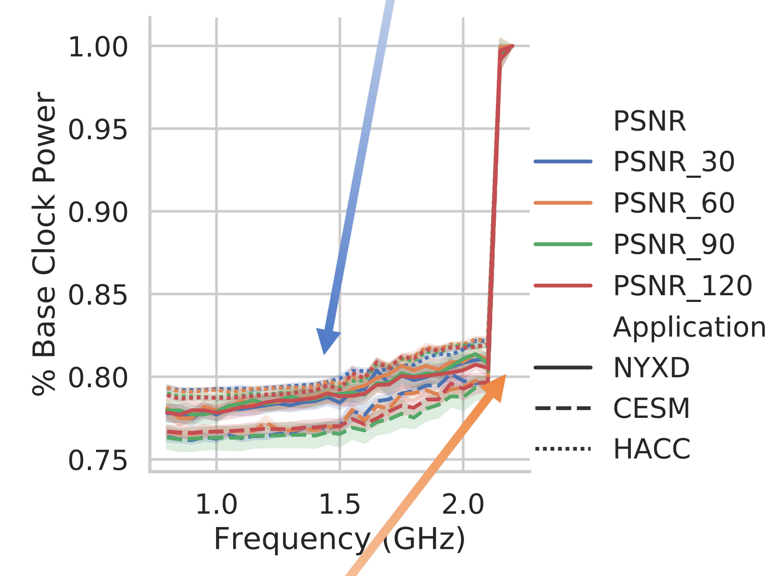

Fig. 4 Power curve to be fit

Small time tradeoff: +5%
Large Power Save: -15%

Fitting the Theoretical Model

Model Fitting:
We fit a model for $P_{SZ}(f, P_{base})$ a function that models power usage in terms of CPU frequency and a constant, P_{base}, base power.

Model	Model Equation P_{SZ}	Sum of Squared Error	Root-Mean Squared Error
Linear	$P_{base}(af + b)$	0.0451	8.488
Quadratic	$P_{base}(af^2 + bf + c)$	0.03694	5.694
Cubic	$P_{base}(af^3 + bf^2 + cf + d)$	0.03062	3.911
Exponential	$P_{base}(ae^{bf})$	0.04475	8.36
Power	$P_{base}(af^b + c)$	0.0229	2.188

Minimizing SSE and RMSE led to selecting Power Model.

Table 2. Model fitting for SZ power function

$$P_{SZ}(f, P_{base}) \approx P_{base}(af^b + c)$$

$a, b,$ and c in Table 3.

Eqn. 2 Form of best-fit black-box model

Parameter	Value
a	$2.237E-9$
b	23.44
c	0.7833

Table 3. Eqn. 2 fit parameters

Applications of Energy Processing

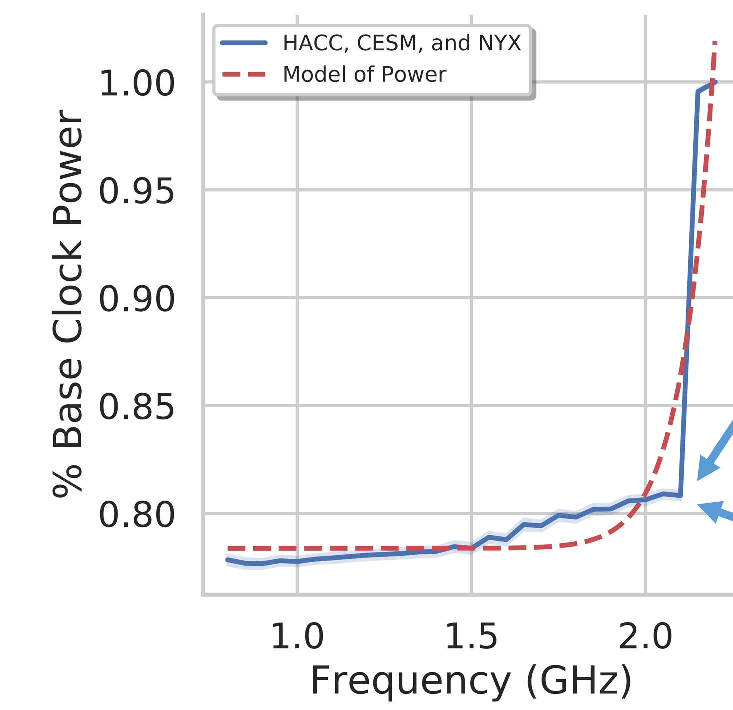

The increase in power usage ~2.1GHz is consistent across diverse experiments.

Model consistent with large spike in power consumption

Fig. 5 Black-box model for the power of SZ

Metric	Power Model Statistic
R^2	0.8296
Adjusted R^2	0.8295
SSE	2.188
RMSE	0.0229

Table 4 relates the goodness of fit for the black-box model.

Table 4. Goodness of fit metrics, relating the model from Fig. 6

Conclusions and Future Studies

Conclusions:
We conclude that power usage can be reduced by ~15% with a 5% time tradeoff.

The predictive model produced accurately depicts power usage for SZ compression and decompression, using CPU frequency.

Future Study Goals:
- Expand to other compressors (e.g. ZFP).
- Test on other CPU architectures.
- Improve model accuracy by quantizing error bound and compressibility.

References

[1] Abhishek Jaiantilal, Yifei Jiang, and Shivakant Mishra. 2010. Modeling CPU Energy Consumption for Energy Efficient Scheduling. InProceedings of the 1st Workshop on Green Computing(Bangalore, India)(GCM '10). Association for Computing Machinery, New York, NY, USA, 10–15. https://doi.org/10.1145/1925013.1925015
[2] X. Liang, S. Di, D. Tao, S. Li, S. Li, H. Guo, Z. Chen, and F. Cappello. 2018. Error-Controlled Lossy Compression Optimized for High Compression Ratios of Scientific Datasets. In 2018 IEEE International Conference on Big Data (Big Data). 438–447
[3] Dmitry Duplyakin, Robert Ricci, Aleksander Maricq, Gary Wong, Jonathon Duerig, Eric Eide, Leigh Stoller, Mike Hibler, David Johnson, Kirk Webb, Aditya Akella, Kuangching Wang, Glenn Ricart, Larry Landweber, Chip Elliott, Michael Zink, Emmanuel Cecchet, Snigdhaswin Kar,and Prabodh Mishra. 2019. The Design and Operation of CloudLab. In Proceedings of the USENIX Annual Technical Conference (ATC). 1-14. https://www.flux.utah.edu/paper/duplyakin-atc19
[4] Daniel Ellsworth, Tapasya Patki, Martin Schulz, Barry Rountree, and Allen Malony. 2016. A Unified Platform for Exploring Power Management Strategies. In Proceedings of the 4th International Workshop on Energy Efficient Supercomputing (Salt Lake City, Utah) (E2SC '16). IEEE Press, 24–30
[5] Salman Habib, Vitali Morozov, Nicholas Frontiere, Hal Finkel, Adrian Pope, Katrin Heitmann, Kalyan Kumaran, Venkatram Vishwanath, Tom Peterka, Joe Insley, et al. 2016. HACC: extreme scaling and performance across diverse architectures. Commun. ACM 60, 1 (2016), 97–104.
[6] JE Kay, C Deser, A Phillips, A Mai, C Hannay, G Strand, JM Arblaster, SC Bates, G Danabasoglu, J Edwards, et al.2015. The Community Earth System Model (CESM), large ensemble project: A community resource for studying climate change in the presence of internal climate variability. Bulletin of the American Meteorological Society96, 8 (2015), 1333–1349.
[7] Ann S. Almgren, John B. Bell, Mike J. Lijewski, Zarija Lukić, and Ethan Van Andel. 2013. Nyx: A MASSIVELY PARALLEL AMR CODE FOR COMPUTATIONAL COSMOLOGY. The Astrophysical Journal 765, 1 (feb 2013), 39. https://doi.org/10.1088/0004-637x/765/1/39
[8] Scientific Data Reduction Benchmark. [n.d.]. https://sdrbench.github.io/. Online. Manuscript submitted to ACM
[9] P. Lindstrom. 2014. Fixed-Rate Compressed Floating-Point Arrays. IEEE Transactions on Visualization and Computer Graphics 20, 12 (2014), 2674–2683.

This material is based upon work supported by the National Science Foundation under Grant No. SHF-1910197.