Computers as Information Channels: A Common Yardstick for the Post-Moore Hardware Zoo

Max Hawkins
mhawkins60@gatech.edu
Georgia Institute of Technology
Atlanta, Georgia, USA

Rich Vuduc (Advisor)
richie@cc.gatech.edu
Georgia Institute of Technology
Atlanta, Georgia, USA

Keywords

information theory, HPC, flop/s, performance measurement, emulation

ACM Reference Format:
Max Hawkins and Rich Vuduc (Advisor). 2026. Computers as Information Channels: A Common Yardstick for the Post-Moore Hardware Zoo. In *Proceedings of The International Conference on High Performance Computing, Networking, Storage, and Analysis (SC '26)*. ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 Summary

We suggest viewing computation as a generalization of a communication channel. Rather than simply copying data, a computer encodes information into data types, performs complex operations (beyond the identity), then decodes the outputs. In his 1948 paper, Shannon abstracted communication away from message semantics [6]. He argued that "semantic aspects of communication are irrelevant to the engineering problem. The significant aspect is that the actual message is one selected from a set of possible messages." We extend this idea to computation.

A communication channel that always emits the same symbol regardless of its input does not transmit any information; likewise, any computation that maps every input to a fixed bit pattern is a "no-op" and is typically optimized away. In contrast, transforming many possible inputs into many distinct outputs (such as arithmetic on encoded real numbers) is closer to our intuition about what represents meaningful computational work.

Three quantities follow. Shannon entropy, $H(X)$, measures the average uncertainty of a random variable or channel inputs/outputs (like bit strings of a two-input adder). Mutual information, $I(X;Y) = H(X) - H(X|Y)$, measures how much of that information a computation transforms and retains. This returns our base performance unit to bits[1]. Compute-channel capacity, $C = maxI(X;Y)$, the mutual information maximized over all input distributions, gives an application-agnostic, hardware-level limit in bits per operation. Multiplying by the operation rate yields bit/s, exactly as flop count becomes flop/s, and independent channels' performance sum by the additivity of mutual information.

1.1 Generalization across Data Types, Bit Widths, Sparsity, and Noise

Mutual information obviates the need to distinguish between flops, integer-ops, or other data types. A unit based on quantities of information is invariant to how the binary strings are encoded or named. They are all just bits.

More importantly, compute channel-capacity accounts for the bit widths of data types. By contrast, simply counting flops means the number of FP64 operations would equal the number of 8-bit equivalents when swapping one type for the other, requiring additional qualification ("FP64 flops" versus "8-bit flops"). Our approach is more in line with the U.S. government's export control approach with Total Processing Performance (TPP) [7].

In our metric, the maximum mutual information of any compute operation is bounded by the smaller of its total input entropy and its total output entropy. Concretely, if an operation takes inputs of bit-widths $w_{\text{in},i}$ and produces outputs of bit-widths $w_{\text{out},j}$, then

$$I_{\max} \approx \min\left(\sum_i w_{\text{in},i},\ \sum_j w_{\text{out},j}\right). \tag{1}$$

As a rough approximation for uniform input distributions, this reduces to TPP's input bit width accounting scheme. When multiplied by the operation rate, channel capacity (in bits per operation) thus provides a first-principles computing performance metric in line with a current government standard.

1.1.1 Sparsity. In traditional flop-based accounting, sparse operations are typically only counted for non-zero inputs. With our information-theoretic perspective, sparsity can be interpreted in terms of entropy. A known zero entry (in structured sparsity or via a mask) carries zero entropy and requires no processing in an ideal case. This makes intuitive sense: if the outcome is certain, there is no information gain and thus no credit for performing the computation. In contrast, an incidental zero value among the entire distribution carries information based on its log probability.

1.1.2 Noise. When noise affects a computation, the mutual information between inputs and outputs decreases. This loss is quantified in our framework. A figure on our poster shows how random bit flips reduce the channel's structural fidelity, decreasing computational throughput.

1.1.3 Emulation. Ozaki-style schemes [3–5] reconstruct an FP64 product from a configurable number of smaller datatype slices (S), so the machine issues a variable amount of sub-datatype tensor-core operations, while the information delivered per result remains

[1] Although the standard unit of entropy and information is technically Shannons [2].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SC '26, Chicago, IL
© 2026 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-x-xxxx-xxxx-x/YYYY/MM
https://doi.org/10.1145/nnnnnnn.nnnnnnn

bounded by what the FP64 output can hold. Counting the low-precision operations therefore rewards a pipeline for the redundancy it introduces, yet treating all emulated operations the same ignores the precision/range/signedness trade-offs inherent in emulation schemes. Bits of mutual information per emulated operation accounts for this and places native and emulated pipelines on a single axis (as shown in a figure on the poster).

1.2 Conclusion

This perspective offers a unifying framework to evaluate the performance of computing systems based on their fundamental ability to encode, manipulate, and retain information. A preprint containing some of these ideas is available on the arXiv [1].

References

[1] Max Hawkins and Richard Vuduc. 2025. Back to Bits: Extending Shannon's communication performance framework to computing. https://doi.org/10.48550/arXiv.2508.05621. doi:10.48550/arXiv.2508.05621

[2] International Electrotechnical Commission. 2025. IEC 80000-13:2025. https://www.iso.org/standard/87648.html

[3] Hiroyuki Ootomo, Katsuhisa Ozaki, and Rio Yokota. 2024. DGEMM on integer matrix multiplication unit. *The International Journal of High Performance Computing Applications* 38, 4 (July 2024), 297–313. https://doi.org/10.1177/10943420241239588. doi:10.1177/10943420241239588

[4] Katsuhisa Ozaki, Takeshi Ogita, Shin'ichi Oishi, and Siegfried M. Rump. 2013. Generalization of error-free transformation for matrix multiplication and its application. *Nonlinear Theory and Its Applications, IEICE* 4, 1 (2013), 2–11. doi:10.1587/nolta.4.2

[5] Katsuhisa Ozaki, Yuki Uchino, and Toshiyuki Imamura. 2025. Ozaki Scheme II: A GEMM-oriented emulation of floating-point matrix multiplication using an integer modular technique. arXiv:2504.08009 [cs.MS]. doi:10.48550/arXiv.2504.08009

[6] C. E. Shannon. 1948. A mathematical theory of communication. *The Bell System Technical Journal* 27, 3 (July 1948), 379–423. https://doi.org/10.1002/j.1538-7305.1948.tb01338.x. doi:10.1002/j.1538-7305.1948.tb01338.x

[7] U.S. Department of Commerce. 2023. Implementation of Additional Export Controls: Certain Advanced Computing Items; Supercomputer and Semiconductor End Use; Updates and Corrections. https://www.federalregister.gov/documents/2023/10/25/2023-23055/implementation-of-additional-export-controls-certain-advanced-computing-items-supercomputer-and

Received 8 August 2026