

Understanding "Information-Friction" for Minimizing Power Consumed in Data Centers and BigData Processing

Pulkit Grover
Carnegie Mellon University
pulkit@cmu.edu

Abstract—This paper discusses a new model of energy consumed in information-processing that can help analyze tradeoffs in energy consumed locally and globally in massive parallel computations. The model, called the "Information-Friction Model," captures the energy required to move information. The model is inspired by the classical Newtonian friction model: just as there are frictional losses associated with moving masses on a surface, there are frictional losses associated with moving information on a computational substrate (be it massive amount of communication between servers or on-chip communication). This circuit-inspired model is more physical (and hence more intimately connected to energy consumption) than the traditional computation models (namely, the Turing machine model and Yao's communication-complexity model), and easier to analyze than earlier VLSI-inspired models. Using this model, we argue that fundamental changes are required in design and implementation of networks in BigData processing systems, including data-centers, to minimize energy consumption. For many problems, we provide both fundamental limits on computation and communication energy, as well as strategies that approach these fundamental limits.

I. INTRODUCTION



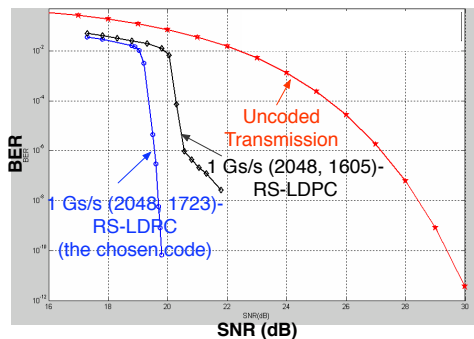


Fig. 1. Just as there is energy associated with moving masses on a surface, there is energy associated with moving information on a substrate (e.g. metal wires on silicon). The "information-friction" model, depicted here and proposed by the author in [1] captures this energy in a relevant yet analyzable fashion.

The works that led to the foundations of the fields of computer science and communications, namely, Turing's work on computability and the Turing machine [2], and Shannon's work on information theory [3], led to revolutions in these fields and are in many ways responsible for the digital revolution that affects our lives so deeply. Nevertheless, modern communication systems and computation systems — and our major concerns with these systems — have come a long way from those that motivated these beautiful models and works. One particular disconnect has been energy: Turing was motivated by the notion of computability and complexity, but not directly with *energy*. For Shannon,



A Shannon-theoretic view of code-performance focuses on just the transmit power



This figure, taken from presentations used in 10 Gbase-T standardization [Rao et al.], shows that the code-choice was guided by principles in traditional information theory.

Fig. 2. The figure shows how traditional transmit-power-focused Shannon-theoretic view was used in code-choice in 10 GBASE-T standardization. The figure on the top is the famous Shannon "waterfall" curve [4], and the one on the bottom is from [5]. Our work has shown that this view can be misleading when choosing codes that minimize transmit + encoding/decoding/processing power [4]. The message is: ECC-choice must depend on, or adapt to, the communication distance.

energy was a concern, but the energy of computation required at the transmitting and receiving ends did not matter. The communication devices of the time were the *telephone*, the *television*, the *telegraph*, etc., (*tele*: greek word that means "far off") which all communicate at long distances: distances where over-the-air transmit power dominates any power consumed in the processing circuitry. But the focus on energy has changed the game.

Let's take a concrete example. The strategies for communi-

cation in data-centers have been chosen based on how close they operate to the Shannon limit (see Fig. 2). Recent theoretical results [1], [4], [6], [7] and empirical observations [8], [9] show that in such situations, operating close to the traditional information-theoretic limits (Shannon’s “channel capacity”) — which minimize just the transmit energy — can be highly inefficient from a system-level (*i.e.*, total) energy-perspective.

Does this energy matter on a global scale? Data-centers themselves consume about 2% of the world electricity, larger than the total electrical energy consumption of most countries. Within a data-center, 15-20% of power (and, in some cases, even up to 50% [10]) is consumed in networking [11]. Thus, even from a world-energy standpoint, networking in data-centers is a non-negligible fraction of the total electrical energy consumption.

Our results have shown that the joint transmit + circuit energy minimization [9] can reduce energy consumption in communication systems significantly at short distances. But can we say something about computation more broadly? Indeed, our recent results (in progress) show that redesign of computational networks for the problem of *sorting* a set of elements could reduce energy consumption significantly. Since problems of this nature occur significantly in BigData computation, there’s hope that this approach can help understand how low power consumption in massive computation can be.

ACKNOWLEDGMENTS

We acknowledge the generous support of NSF-ECCS-1343324, seed grants from the NSF Center for Science of

Information NSF-CCF-0939370, and a startup grant from Carnegie Mellon University.

REFERENCES

- [1] P. Grover, ““Information-friction” and its impact on minimum energy per communicated bit,” in *IEEE International Symposium on Information Theory (ISIT)*, Istanbul, Turkey, Jul. 2013.
- [2] A. M. Turing, “On computable numbers, with an application to the entscheidungsproblem,” *Proceedings of the London mathematical society*, vol. 42, no. 2, pp. 230–265, 1936.
- [3] C. E. Shannon, “A mathematical theory of communication,” *Bell System Technical Journal*, vol. 27, pp. 379–423, 623–656, Jul./Oct. 1948.
- [4] P. Grover, K. Woyach, and A. Sahai, “Towards a communication-theoretic understanding of system-level power consumption,” *IEEE Journal on Selected Areas in Communication*, Sep. 2011.
- [5] S. Rao, R. Hormis, and E. Krouk, “The 4D-PAM8 Proposal for 10GBASE-T,” Nov. 2003, P802.3 Task Force Meeting, http://www.ieee802.org/3/10GBT/public/nov03/rao_1_1103.pdf.
- [6] P. Grover, A. Goldsmith, and A. Sahai, “Fundamental limits on the power consumption of encoding and decoding,” in *Proc. IEEE International Symposium on Information Theory (ISIT)*, Jul. 2012.
- [7] P. Grover, A. Goldsmith, A. Sahai, and J. Rabaey, “Information theory meets circuit design: Why capacity-approaching codes require more circuit area and power,” in *Proceedings of the Allerton Conference on Communication, Control, and Computing*, Monticello, IL, Sep. 2011.
- [8] K. Ganesan, P. Grover, and A. Goldsmith, “How far are ldpc codes from fundamental limits on total power consumption?” in *Proceedings of the Allerton Conference on Communication, Control, and Computing*, Monticello, IL, Oct. 2012.
- [9] K. Ganesan, P. Grover, and J. M. Rabaey, “The power cost of overdesigning codes,” in *IEEE Workshop on Signal Processing Systems (SiPS)*, Oct. 2011.
- [10] B. Heller, S. Seetharaman, P. Mahadevan, Y. Yiakoumis, P. Sharma, S. Banerjee, and N. McKeown, “Elastictree: Saving energy in data center networks,” in *NSDI*, vol. 3, 2010, pp. 19–21.
- [11] A. Greenberg, J. Hamilton, D. A. Maltz, and P. Patel, “The cost of a cloud: research problems in data center networks,” *ACM SIGCOMM Computer Communication Review*, vol. 39, no. 1, pp. 68–73, 2008.