# Reducing Energy Consumption and Decentralizing Computing through Heat Redistribution

Lincoln S. Stevens, Mackenzie J. Wibbels, Valerie Wilkinson, and Kenneth S. Stevens
Granite Mountain Technologies, Inc.

***Abstract*—A byproduct of computations performed by semiconductor devices is heat. Many compute intensive functions are currently performed in the cloud by hardware placed in data centers, where heat must be removed at additional energy cost. Granite Mountain Technologies is exploring an energy conserving model and technology for decentralizing computation from data centers to areas where heat is required. This compute model can provide vast energy savings benefits, while also redistributing monetary gains from large organizations capable of running data centers to individuals. We report on a decentralized computing appliance implemented with relative timed integrated circuits in the 12nm technology node. The appliance can accept calculations at a variable rate which allows heat created from computations to achieve the desired temperature of a room. The benefits of applying this approach to residential heating is evaluated.**

## I. Introduction

As semiconductors perform computations, they produce heat. Heat is normally an undesired byproduct of computation, and requires additional energy and expense for removal. In a centralized model, computations are performed in large data centers, where additional energy is expended to remove unwanted heat. Compensation for computation benefits the centralized few.

A decentralized model moves compute power to the edge. Strategically placing edge computations where heat is needed reduces overall energy consumption in two ways. First, the heat generated as a byproduct of computation is relocated to environments where heat is needed, replacing energy usage that would otherwise be spent in heat generation. Second, the energy required to cool data centers is eliminated. Energy savings are substantial. An additional societal benefit is that profit from computations can be more broadly distributed.

There are various commercial and consumer applications for this distributed model, including agricultural, manufacturing, and heating living and working spaces in cold environments. While each application may have different specific nuances, following are the general requirements for the model. (a) The ability to distribute computations across the network. (b) The ability to vary the work load rate at remote nodes. (c) Relatively low network bandwidth and loose latency requirements. (d) A reasonable reward structure for the computations. (e) Hardware with high reliability, robustness, and lifetime. (f) Reasonably low hardware cost, or an incentive to move expensive hardware from the data center to the edge.

Decentralizing computations into residential areas is feasible and beneficial. Data centers in the United States consumed approximately 70 B kWh per year in 2014 and accounted for approximately 2% of the total electric use [1]. Cooling accounted for half of the total energy consumption in data centers in 2007 [2], which is no longer needed when heat is a primary application of the computation. Residential energy cost accounted for 21.8% of US energy consumption in 2014, of which 41.5% was spent on heating spaces [3]. In 2015 60% of homes used heating applications which burn hydrocarbons, a non-renewable natural resource which produces green house gasses [4]. Every dollar of computation which performs heat redistribution into homes replaces the equivalent energy use, usually produced by burning hydrocarbons such as natural gas or propane. Repurposing a fraction of the energy used by semiconductors can significantly reduce hydrocarbon use and save money.

From the consumer perspective, there are four key considerations for compute generating heat appliances: heat, noise, profit, and cost. Because these are semiconductor devices, the up front cost can be substantial, similar to solar panels. If the device is too costly to purchase, or too loud for occupants, the devices will not be used.

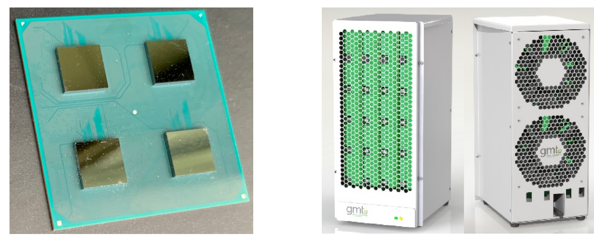

(a) (b)

Fig. 1: The die and substrate (a), and complete system (b).

## II. Implementation Overview

A Bitcoin "mining" appliance was selected for the heat redistribution prototype. Bitcoin mining is an ideal candidate for heat redistribution because it exhibits attributes (a)–(f) listed above. The current worldwide electricity use for Bitcoin mining is about the same as the energy usage of data centers in the United States in 2014.

Bitcoin is a virtual currency used in a global financial network called the Bitcoin network [5]. Like other financial applications, transactions are validated and recorded in a ledger. A unique aspect of the Bitcoin network is that the ledger is both public and secure, and transactions are fully automated.

Ledger transactions are based on the one way SHA-256 hash function. This poses significant power delivery and thermal management challenges in chip design because pipeline and data activity factors are near 100% and 50% respectively. High performance modes of our appliance require up to 1,000 amps per printed circuit board. Heat removal to the environment is key for this application.

The relative timing design flow was used to implement a mining ASIC [6]. Relative timing supports modular, composable hardware implementations, with rich, deep concurrency, and significant power-performance-area benefits over traditional design approaches. The chip (Fig. 1(a)) was fabricated in TSMC's N12FFC technology node.

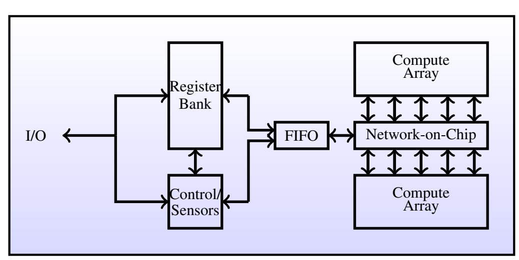


Fig. 2: The B-1 System-on-Chip Chip Block Diagram

The advantages of using relative timing (RT) are evident in multiple levels of the appliance. Chip power, energy, and area are reduced due to relative timing optimizations. The modular nature of RT design greatly reduces design and implementation time and cost. Each engine in the chip operates at a unique frequency based on local temperature, process, and voltage conditions increasing performance and efficiency. At the system level, performance and power are easily controlled by simply adjusting the supply voltage to the chips.

The heating appliance is shown in Fig. 1(b). Each appliance contains four printed circuit boards with 10 packages. Each package contains four chips, and each chip contains 250 compute engines. The system was designed with the ability to arbitrarily scale the voltage delivered to the integrated circuits (ICs) from 470 mV to 320 mV. This allows the heat requirements to be directly controlled by the voltage supplied to the ICs. By lowering the voltage, heat is quadratically reduced, performance linearly degrades, and computation energy efficiency is linearly improved. The energy efficiency of GMT's chip is similar to Bitmain's Antminer S9 designed in a similar technology node. The architecture of the chip is shown in Fig. 2.

Considering only electricity cost, income $i$ for Bitcoin mining is calculated as follows. Let $f$ be the average mining rate in blocks per hour, $b$ the mining payout in Bitcoin per block, $r$ the exchange rate in dollars per Bitcoin, $n$ is the network hash rate in TH/s, $p$ power cost in \$/kWh, and the appliance efficiency $e$ in kW/TH/s. Eqn. 1 shows income $i$ in \$/h/TH/s. The break even boundary is plotted by efficiency $e$ in Fig. 3 given today's conditions where $f = 6$, $b = 12.5$, $r = 8700$, and $n = 105 \times 10^6$ TH/s.

$$i = \frac{fbr}{n} - pe \tag{1}$$

Mining income based on power cost and chip efficiency can be observed based on the graph in Fig. 3. Income is earned when power $p$ costs less than the break even boundary for a given appliance efficiency. An income of 2.4¢/kWh and 7.4¢/kWh is achieved at electricity costs of 10¢ and 5¢ per kWh at an efficiency of 50 W/TH/s. Voltage scaling linearly changes efficiency from $e$ to $e_v$, producing income in \$/kWh as $i_v/e_v \times e_v/e$. This is plotted as red and blue dotted lines in Fig. 3 for power costing 10¢ and 5¢ respectively across the yellow efficiency range from 50 to 100 W/TH/s.

When energy used for computation would otherwise be used for heat, $p = 0$ in Eqn. 1. This makes savings $s$ equivalent to mining revenue $s = fbr/n$. Savings can be plotted in terms of \$/kWh dividing

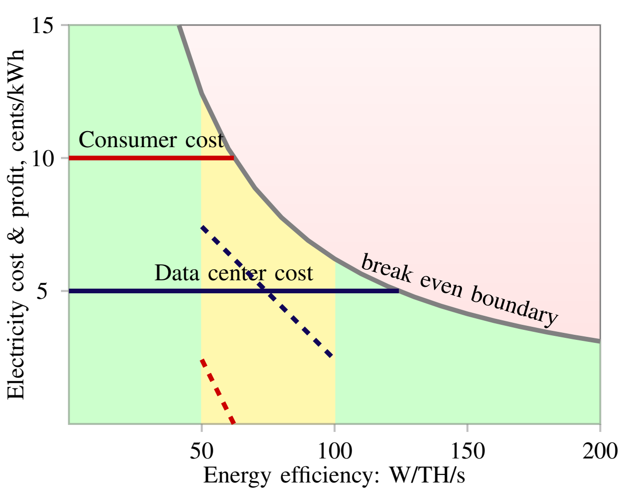


Fig. 3: Profit Calculations

s by e, making it equivalent to the break even line in Fig. 3. Therefore under the heat redistribution model, there is *always* energy savings and income earned, but at increased efficiency higher savings accrue.

## III. Results

To test the heat redistribution model, appliances were placed in a 1400 square foot mountain dwelling where winter temperatures routinely drop below 0°F. Appliances were used for heat assistance through the colder months. Noise levels ranged from 40–70 decibels, compared to 90–100 for the Antminer. Propane usage dropped to 60% of the previous year (by 425 gallons), whereas electric usage increased by 61% (5874 kW hours) as shown in Fig. 4. Overall heating costs were similar, as the gas bill dropped by about \$712 and the electric bill raised by approximately \$634. The appliance earned 0.1 Bitcoins, which is worth approximately \$870 today.

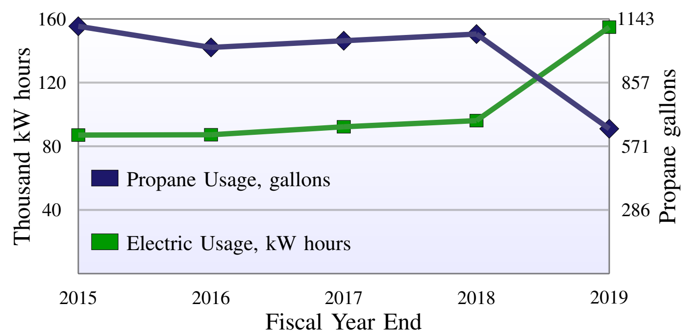


Fig. 4: Energy usage, heat repurposing in 2019

## IV. Implications and Conclusion

The project confirmed that a new model for edge computation that redistributes heat can be achieved through decentralization of computing power coupled with the ability to control heat output by varying the voltage delivered to a processor. An appliance was built with custom ICs designed using relative timing in the 12nm technology node with a focus on thermal management. The appliance was tested by heating a mountain home, resulting in an aggregate reduction in heating costs of \$948, and reducing propane usage by 425 gallons. The ambient noise could be sufficiently masked to allow the appliance to be continually operated.

If this type of compute model is widely used, millions of gallons of hydrocarbons would not be burned, and billions of dollars could be saved by households across the world, making the world a greener place for the environment and for adopters.

## References

[1] A. Shehabi, S. Smith, D. Sartor, R. Brown, M. Herrlin, J. Koomey, E. Masanet, N. Horner, I. Azevedo, and W. Lintner, "United States Data Center Energy Usage Report," Ernest Orlando Lawrence Berkeley National Laboratory, Tech. Rep. LBNL-1005775, June 2016.

[2] R. Brown, E. Masanet, B. Nordman, B. Tschudi, A. Shehabi, J. Stanley, J. Koomey, D. Sartor, P. Chan, J. Loper, S. Capana, B. Hedman, R. Duff, E. Haines, D. Sass, and A. Fanara, "Report to congress on server and data center energy efficiency: Public law 109-431," Lawrence Berkeley National Laboratory, Tech. Rep. LBLN-XXXXX, August 2007.

[3] Energy data facts. U.S. Department of Energy. [Online]. Available: https://rpsc.energy.gov/energy-data-facts

[4] Residential Energy Consumption Survey (RECS). U.S. Energy Information Administration. [Online]. Available: https://www.eia.gov/consumption/residential/data/2015hc/php/hc1.1.php

[5] (2008, October) Bitcoin: A Peer-to-Peer Electronic Cash System. Bitcoin.org. [Online]. Available: https://bitcoin.org/bitcoin.pdf

[6] K. S. Stevens, R. Ginosar, and S. Rotem, "Relative Timing," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 1, no. 11, pp. 129–140, Feb. 2003.