

An Analysis of the Increasing Processing Power of Modern Microprocessors

Noah Kenney

December 20, 2022

1 Background Information

Traditional computational power calculations rely on the assumption that additional processing power is best achieved by a quantitative increase in transistors, and particularly through a quantitative increase in the number of metal-oxide-semiconductor field-effect transistors (MOSFETs). However, advancements in microprocessor architecture could fundamentally alter the method in which additional processing power is achieved, thus rendering traditional computational power calculations meritless.

Modern technologies such as Artificial Intelligence, deep learning, and neural networks all require a high level of processing power with minimal demands on resources (particularly energy). Thus, it is worthwhile to consider the ability of various chip structures to provide an increase in computational processing power without further increasing strain on resource consumption.

2 Moore's Law

To begin this analysis, it is necessary to understand the current expected rate of increase in processing power of standard microprocessors. Moore's Law is able to provide such an estimation of the rate of increase of processing power relative to cost over time, given a variable number of biennial periods. Thus, determining the cost of processing power relative to a fixed computational unit can be accomplished by deriving an equation that serves as a mathematical representation of Moore's Law.

$$x = \frac{(2c)^{(n-1)}}{2a}$$

Thus:

(x) is equivalent to the cost per quantitative unit of computational power, calculated in quantity of transistors given a fixed square unit size.

Where:

($2c$) represents current processing power in transistors

(n) represents the number of biennial periods that have passed since the base year referenced in (c)

(2) represents the cost of computational power being cut in half on a biennial basis

(a) represents the number of square units in the computer chip the calculation is being performed on

While cost is not necessary to this analysis of processing power in microprocessors, it does serve to show the relative reliability that can be placed upon Moore's law. In recent years, many have criticized Moore's Law for showing diminishing accuracy as the number of transistors fails to double in density on a biennial basis. Its presumption of cost estimation also shows diminishing accuracy, particularly in light of recent chip shortages, which have increased chip costs.

For reference, Moore's Law would assume that microprocessors in 2023 would have approximately 114 billion MOSFETs, given that microprocessors in 2021 have as many as 57 billion MOSFETs. Additionally, we would expect a MOSFET scaling of approximately 2.5-3 nm. Concurrently, Moore's Law would assert that the cost of the given microchip would be cut in half over the biennial period between 2021 and 2023. However, currently available microprocessor schematics do not indicate that such results will be achieved, with the potential exception of MOSFET scaling (2.5-3 nm).

3 Microprocessor Advancements

We must also consider research proposing that Moore's Law, and traditional calculations of processing power relative to cost, lack merit given the advent of microprocessors with vastly different chip structures. As an example, neuromorphic computing utilizes a vastly different microprocessor structure, relying on the use of a neuro-biological architecture that represents the neurons in the human brain through the use of very-large-scale integration (VLSI) systems and electronic analog circuits.

Additionally, microprocessors currently rely on the use of nanotechnology, which allows for a limited range in dimension and tolerance from a scale of 1-100 nanometers. However, current rates of MOSFET scaling indicate that the necessity of range in dimension and tolerance smaller than 1 nm will soon exist. As a result, work is already being done on mechanical manipulation of atoms that achieve a scale smaller than nanotechnology (scale of 10^{-9}). Examples of such technology branches include picotechnology (scale of 10^{-12}) and femtotechnology (scale of 10^{-15}). Both are currently considered hypothetical, but significant

progress has been made toward attaining feasibility. Thus, processing power would increase significantly in a short amount of time, but would likely lack the ability to achieve the exponential growth of transistors that Moore's Law predicted based on microprocessors built using nanotechnology.

4 Power consumption of Processors

Traditional computational processor power dissipation utilizes resistance to consume power, eventually converting it to heat. Naturally, as processing power increases, the production of heat also increases. Thus, increasing processing power relies on an assumption of fundamental increase in thermal heat regulation and higher tolerance in computer resistors.

We can model CPU power consumption using the following equation:

$$P_{cpu} = P_{dyn} + P_{sc} + P_{leak}$$

Most relevant when evaluating constraints of microprocessors, P_{leak} represents the quantifiable power lost from transistor leakage. It is reasonable to figure that there are possible, if not probable, constraints on modern microprocessors resulting from a lack of improvement in dynamic and/or short-circuit power consumption. At best, it is a driver of processing power costs. Reducing transistor leakage is one possible approach to overcoming this challenge. Thermal throttling has historically been utilized to prevent the overheating of chips and optimize the processing speed, in conjunction with DRAM or, in rare cases, SRAM. However, thermal throttling defeats the purpose of higher theoretical transistor counts, except perhaps in the case of highly specialized computational systems. It is further possible that alternative types of processors will rely on an entirely unique thermal regulation system that renders discussions of heat regulation and transistor leakage futile.

5 Alternative Types of Processors

The development and increased usage of unified memory chips in modern computers also allows for a greater allocation of processing power without increasing transistor count. This is provided by the single chip architecture, but still uses silicon chips. There are several alternatives to silicon chips promising to deliver the same, or higher, processing power without the traditional chip architecture.

Quantum computers, largely rooted in quantum physics instead of computer science and engineering, uses qubits to provide processing power. They are highly efficient for large computational processes as they use statistical calculations to compute probability of the state of the particle (1 or 0). By utilizing subatomic particles (generally electrons or photons), the particles can exist in

both states (1 and 0) at the same time, unlike silicon competitors. It remains to be seen whether quantum computers are equally effective and efficient when used for everyday computing tasks. The cost of developing qubits also prohibits quantum computers from reaching market scale currently.

A second alternative to silicon computing is graphene and carbon nanotubes, an extremely new technology with significant long term potential. The unique hexagon lattice shape of graphene makes it much stronger than other materials. It is also much smaller (one atom thick), and available in large enough quantities to make it commercially viable. It is highly probable that a new thermal regulation system would need to be developed to support graphene based chips as it cannot switch electrical current on and off as binary powered chips do.

Nanomagnets, another alternative to silicon chips, rely on cellular architecture (making them particularly suited for mobile devices). The state of magnetization determines when (or if) binary information is sent/received. It uses opposite magnetic poles (north and south) to transmit and process binary data instead of electrical current which reduces power consumption down to a negligible amount.

6 Summary

After thorough analysis, it seems reasonable to state that Moore's Law will not continue to provide an accurate method of predicting processing power and computation processing cost for any significant period of time. Traditional silicon based chips fail to add transistors at the rate of speed necessary to line up with the biennial period predictions (with the exception of silicon chips utilizing a unified DRAM approach). At the same time, alternatives to silicon chips become more feasible, particularly given that they do not rely on traditional thermal regulation systems. In the next decade, it is highly probable that these alternatives to silicon computing will fundamentally change the way processing power is measured in modern microprocessors.