Do Voltage and Clock Switching Really Save Energy?

Revised March 13, 2006

Edward Herbert
Canton, CT 06019

The energy consumed to complete a task requiring a fixed number of clock cycles is explored. The instantaneous power is lower at a reduced clock frequency and reduced voltage, but the energy (power x time) is comparable. This suggesst that voltage and clock switching have very marginal benefit if the processor operates only when working. By tuning on a processor core only when an operation occurs and turning off immediately afterwards, all of the hardware, firmware and software associated with voltage and clock switching can be eliminated.

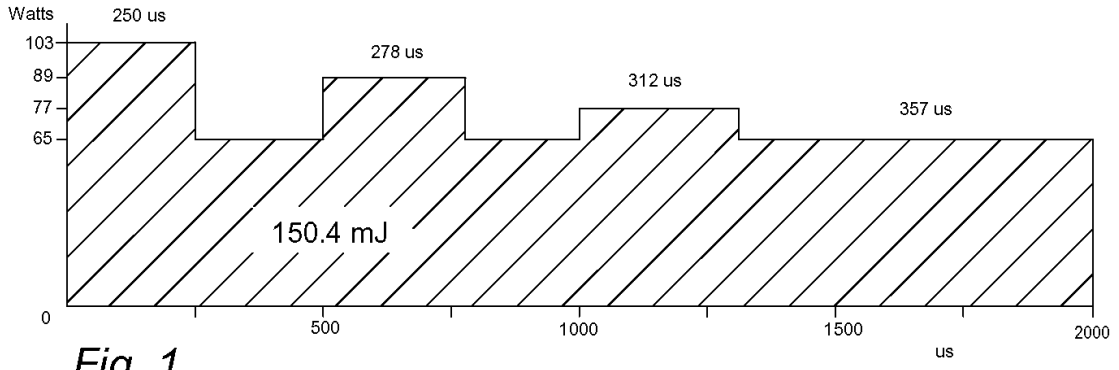

Fig. 1

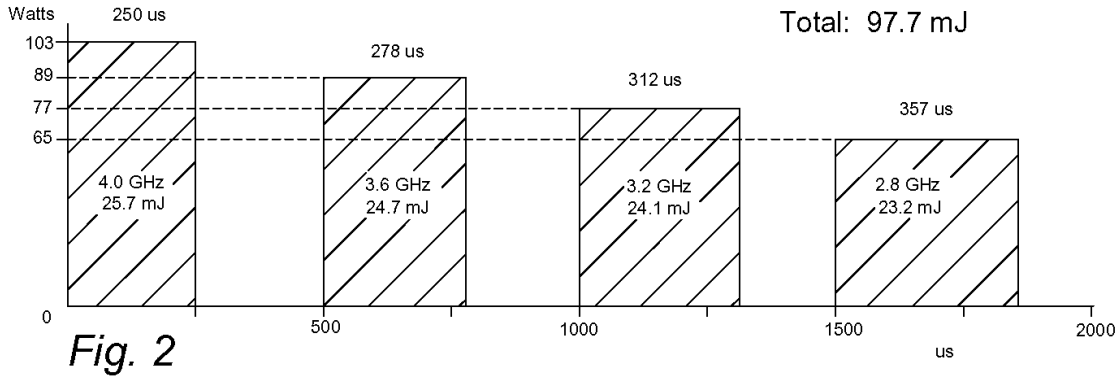

Fig. 2

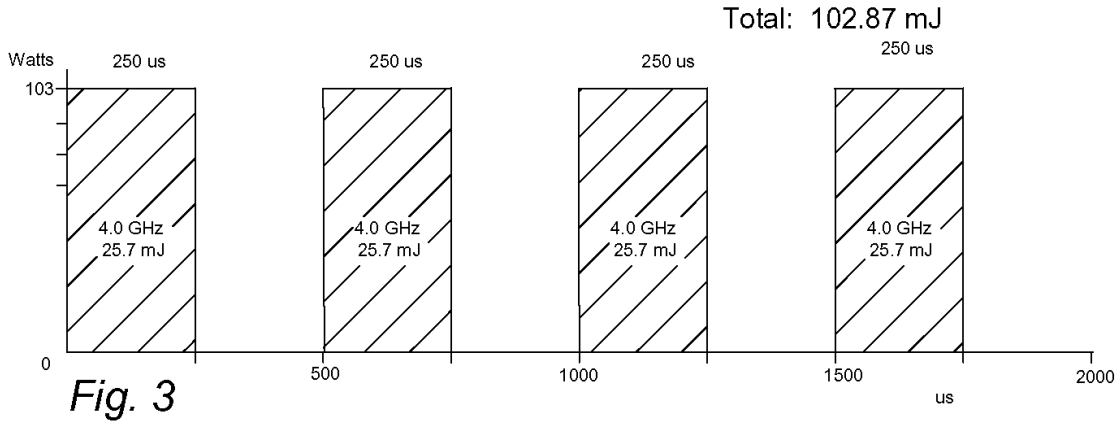

Fig. 3

Figure 1 shows the energy consumption (the area under the curve) of a processor that has a reduced voltage when idle between tasks, and operates at different voltages and clock frequencies.

Figure 2 shows the same processor performing the same tasks and operating ad different voltages and clock frequencies. The processor is turned off between tasks. Turning off between tasks saves significant energy.

Figure 3 shows the same processor performing the same tasks, except always at full voltage and full clock speed. The processor is turned off between tasks.

A processor consumes less power if it is operated at a lower voltage, but to do so, it must also reduce its clock speed. Some processors have very complex reduced power states, such as the AMD Coon'N'Quiet and Intel Foxton, involving up to 64 voltages and 64 clock speeds. A representative chart, from an Intel paper, shows the following power savings at reduced clock frequency and reduced voltage:

4.0 GHz1.39 V	103 W
3.6 GHz1.33 V	89 W
3.2 GHz1.27 V	77 W
2.8 GHz1.20 V	65 W

This is a pretty impressive power savings! Figure 1 shows the energy consumed assuming that the clock and voltage are reduced to the lowest power state except when a demanding task occurs.

However, considering the energy to perform an operation, the answer is quite different. Taking a hypothetical operation that takes 1,000,000 clock cycles, I calculated how much energy the processor would consume to complete the operation at the various clock frequencies and voltages. For example, at 4.0 GHz, 1,000,000 clock cycles takes 250 us, and at 2.8 GHz, the same operation takes 357 us.

The calculation for energy (in joules) is:

energy = power x time

The time to do a task needing a fixed number of clock cycles will obviously take longer if the clock frequency is slower. The energy consumed by a given task requiring 1,000,000 clock cycles can therefore be expressed as

energy = power x (number of clock cycles) / clock frequency.

Solving for the hypothetical 1,000,000 clock cycle task using the above relationship of power to clock frequency gives the results shown in Figure 2, and tabulated below:

4.0 GHz	0.0257 J
3.6 GHz	0.0247 J
3.2 GHz	0.0241 J
2.8 GHz	0.0232 J

This is pretty slight energy savings *per task*, about 10%. This means that if the processor always ran at full bore (maximum voltage, maximum clock), but turned on *only* when needed and turned off *immediately* after a task was done, there would be no significant advantage in energy consumption for voltage and clock switching. This is shown in Figure 3.

All the software control of the power would be unimportant, as would voltage switching and clock switching, other that on-off. This would be a huge simplification of software, firmware and hardware. I bet that there would be some power savings there as well, so, there may be no penalty at all in discarding the complexity of voltage and clock switching and using just a simple on-off control.

Assuming that a low power supervisory processor could handle the pointers and registers, and turn the core on only when there was a priority job to do, the other tasks in firmware and the hardware to do voltage and clock shifting, as well as the software programming, all become unnecessary. They just need the SCPC, with an on-off control and a static VID.

The processor community is very heavily invested in clock and voltage switching, and they may not be easily persuaded that it is not particularly beneficial for energy savings. Yet the proof is there, if energy per task is calculated. This is one area where a power supply that could go from VID, full load to 0 V, 0 A and back in 2 us may save a lot of hardware, firmware and software. *And, the power is zero when the processor is off.*

I have seen figures suggesting that most servers operate below 4 percent duty-cycle for about 85 percent of the time. Visualize figure 3 with 6,250 us between the pulses with 0 W most of the time. If that is true, completely turning off the processor core could save 96 percent of the power over keeping it active in a low voltage mode. The overhead circuits would consume some power, but a net power savings of 90 percent seems realistic.

Interesting links:

Enterprise Platform Power Efficiency for Future Server Systems
Tomm Aldritde, Manager – Enterprise Power & Thermal Technology Labs
Intel Corporation — IBM Platform Technology Symposium September 14-15, 2004
Kahler Grand Hotel, Rochester, MN

Taurus: A Taxonomy of the Actual Utilization of Real UNIX and Windows Servers,
David G Heap, Principal IT Consultant, IBM Enterprise Server Group, Somers, NY, USA