THE EDITOR'S VIEW

Power Issues May Limit Future CPUs

Portable Systems Particularly Vulnerable As Power Dissipation Climbs

$P = cv^2f$. It's a fact of life; every electrical engineer learns it in school. Once in the real world, though, it's easy to forget how simple rules control the big picture. In this case, the factors affecting the power consumption of a microprocessor are getting out of balance, forcing CPU designers to pay more attention to this specification in future products.

The equation says the power consumption of a CMOS circuit depends on its total capacitance *(c)*, supply voltage *(v)*, and clock frequency *(f)*. Of these factors, voltage, raised to the second power, has the biggest effect. To increase performance, CPU designers lean heavily on higher clock speeds, driving up power consumption. To compensate, IC process engineers attempt to reduce the capacitance and voltage.

In a typical process shrink, the clock speed of a microprocessor increases by about 50%. The capacitance of the chip is reduced by about 20%, and the voltage drops by an average of 15%. Combining these factors results in a change of $0.8 \times 0.85^2 \times 1.5$, an average decrease of 13% in power.

Of course, in some cases a process shrink can significantly reduce power. A well-known example is the original 0.8-micron Pentium, which decreased from 16 W (maximum) to 10 W when shrunk to a 0.5-micron process, even though the clock speed increased from 66 to 100 MHz. This was a short-term decrease due to the step-function drop in supply voltage from 5 V to 3.3 V. After two process shrinks and one voltage drop, Pentium's power consumption has climbed back to 17 W for the 200-MHz version, almost exactly where it started.

The problem comes in moving to the next-generation processor. The number of transistors that can be placed on a single die approximately doubles with every process shrink. CPU designers take advantage of these transistors to increase the complexity of their processors, hopefully improving the instruction throughput. These extra transistors, as well as all the metal traces needed to connect them, significantly raising the total capacitance *(c)* of the circuit.

The Pentium Pro CPU chip, for example, has 70% more transistors than Pentium and consumes 50% more power at 200 MHz, the maximum clock speed for both processors in Intel's 0.35-micron process. (Since Pentium Pro is designed to achieve higher clock speeds than Pentium in the same process, the power gap could be even greater.)

To continue to double performance every 18 months, CPU designers must push both transistor count and clock speed to the limits of process technology. Even with decreasing supply voltages, each new process generation raises power by 74%, assuming designers take advantage of a doubling in transistor count. A quick look at history shows the 1.0-micron 486 at 4 W, the 0.8-micron Pentium at 16 W, and the 0.5-micron Pentium Pro at 25 W.

Because power consumed must be dissipated, increasing power leads to hotter chips. For desktop systems, additional emphasis on cooling (bigger fans and/or fans mounted on the CPU heat sink) should be sufficient for the foreseeable future, providing a slight increase in the cost and complexity of systems. Dissipating heat in a portable system, however, is much more difficult. Furthermore, power-hungry CPUs will choke the battery life of portables.

Today, the market is demanding notebook systems that offer the same performance as desktops. To meet the power requirements of portables, however, Intel has been forced to reduce the supply voltage of its Pentium processors, limiting their clock speed. With the increased power demands of the P6 CPU, Intel will need to take a similar approach, but the gap between the fastest notebook and desktop versions is likely to grow, not shrink. Other processor vendors face the same issues; Cyrix, for example, has essentially given up on marketing processors for the notebook market.

All is not lost. The key criterion for portable systems is typical, or average, power consumption, not the maximum. Gated clocks and other design techniques can reduce the typical consumption. These techniques have already been used extensively in Pentium to wedge that chip into the power envelope of notebook systems, so it isn't clear they will provide greater relief in future designs.

More radical alternatives include asynchronous circuits that use power only when doing real work; improved communication between the OS and the CPU so the processor runs only as fast as needed for the task at hand; IC processes that support high clock speeds at very low voltages; and abandoning CMOS altogether in favor of a technology with lower power consumption.

Over the next several years, the key tool for battling rising power dissipation will be reducing the supply voltage. Rapid voltage changes will require flexibility on the part of system designers and interface circuitry. During this period, desktop processors are likely to get hotter, and portable systems will lag further behind in performance. Within 5–10 years, one or more of the radical alternatives noted above are likely to become commonplace, giving microprocessors a chance to cool off. M

Linley Gwennap