

minimizing power via strong process voltage scaling

Contributed by Maciej Bajkowski
Tuesday, 14 August 2007

Not so long ago, the name of the game and the claim to fame was processor speed and megahertz. But as with most technology related things, the landscape changes quickly. These days, it is all about who can do more with less; less power that is. To approach this challenge, companies have to some degree been simplifying their cores, and instead of raising the megahertz count have been increasing the core count. The argument being, that several cores running slower but in parallel are more efficient than one core running at maximum speed, from the throughput and power envelope level point of view. However, it does not take a genius to figure out that if you put enough cores onto a single die, the power envelope per core decreases if the overall power for the chip is to stay the same. Additionally, as you increase the core count, the die area for the chip is likely to increase, which can lead directly to yield problems. A less leaky process and smaller feature sizes are no substitute for sound power management techniques and smart design. Addressing this very issue, Heinrich Hillmayr from Texas Instruments' (TI) German division, posted an interesting article titled *Minimizing Power Consumption at the Chip Level* over at PowerManagementDesignLine.com (This must be one of the longest domain names I've ever come across!). If you are a seasoned circuit designer you can skip the first page of the article which briefly describes the basic relationships between voltage, leakage and dynamic currents, and power. If not, this might be a nice little review especially if you have an interview coming up. The second page is where the article gets interesting. In a perfect world there would be neither variation in a wafer nor any variation between wafers, but semiconductor manufacturing is far from perfect. Thus in some regions, transistors are stronger than the nominal transistor, meaning that at a given voltage they conduct more current, while in other regions they are weaker and thus conduct less current. This causes two problems: Transistors that are located in the strong region will have a higher leakage current thus increasing the overall core power, while circuits in the weak region will have a problem meeting speed targets. The idea therefore is to diagnose the strength of the silicon, and to lower the voltage in strong silicon regions in order to decrease leakage currents while still meeting timing. The article does not disclose whether the silicon strength is determined by some sort of on-die sensors but simply states that the mechanism is based on embedded information. I imagine that designing an on-die sensor to accurately determine transistor strength could be quite challenging, but not impossible. Whether such a design would be cost effective in the end is of course another matter. The article also describes voltage scaling based on temperature and has several charts to explain the details mentioned above and is definitely worth a quick glance.