The Return of Silicon Efficiency

Simon Knowles

The economic and physical forces which have always shaped the business of digital chip design are again evolving to change the priorities of designers. Key physical trends include the end of gate oxide thickness scaling, and the effect of small dopant populations on threshold voltage variance. Key economic trends include the need to tolerate specification shift and design error, and the need to amortise chip development cost over multiple market sockets.

Macroscopically, the pursuit of speed is giving way to new emphasis on power efficiency, and the pursuit of time-to-market is giving way to new emphasis on application flexibility. A defining characteristic of the new order is the trend towards more software-defined functionality, and therefore implicitly to processor-centric hardware platforms.

Together these trends represent a shift away from the “transistors are free” mantra of the ASIC era, back towards a value set more reminiscent of 80’s custom design – squeezing the most from each transistor. Silicon efficiency is important again. Design idioms and EDA tools must change accordingly.

This talk will review some of the economic and physical factors currently shaping digital chip design, and will quantify pivotal trade-offs between speed, power, cost, and design methods.