Preface

Progress in information technology continues at a super-exponential growth rate. This has been true for more than one hundred years and over many different generations of computing technology—mechanical, vacuum tube, transistor, integrated circuit. The progress has come about through improvements of existing technologies as well as the discovery of new ways of doing things. What we now think of as Moore's law, i.e., the tendency for the number of transistors on a chip to double every 18 to 24 months, is just the latest chapter in this story of progress.

In the next ten to twenty years this should generate staggering results. For example, today's personal computers can perform approximately one billion operations per second. For comparison, that is roughly the operational power of the brain of a lizard, or a dinosaur. In twenty to twenty-five years, given a continuation of the super-exponential growth rate, a personal computer will be able to perform about 10^{15} operations per second (a petaflop). By various estimates, that is the operational power of the human brain. In other words, in the next twenty years, the growth in computing power should correspond to hundreds of millions of years of human evolution. These computers of the future will not think as humans do, but the progress will be staggering nonetheless.

The question this issue of the *Journal* addresses is whether such advances will continue. What is the future of information technology? Technologists see many of the substantial barriers to further advancement. In silicon technology, CMOS scaling appears to be approaching an end, while rapid progress in optical lithography seems to be nearly impossible. It is difficult to imagine how we can make the needed insulating layers thin and perfect enough to fabricate the transistor gates of the future, or control the statistical fluctuations in doping materials. Interconnection layers are struggling with growing chip sizes and with speed-of-light limitations. In magnetic storage the superparamagnetic limit seems to be not far away, and recording head scaling a growing problem. Nonetheless, history has shown that we have always found new ways to extend existing technologies or develop new ones to go beyond those which are reaching their limits.

The papers in this issue describe both the evolving technologies and the breakthroughs that may fuel the information age of the future. Evolving magnetic storage technologies are described in the paper by Thompson and Best, together with some of the barriers to progress. Even given these, the authors see possible improvements within variations of magnetic storage technology by factors of up to 100 times. Beyond that, new techniques will probably be required. Two of these are discussed here: holographic storage and probe tip storage. Holographic storage, described by Ashley et al., has the advantage of being a volumetric technique with the potential of huge data rates. One method of probe tip storage, the Millipede, is described in the paper by Vettiger et al. They have demonstrated that this form of storage, based on the operation of arrays of atomic force microscopes (AFMs), is capable of large areal densities. One can even imagine extending the approach into an atom-level storage regime.

In the silicon arena, the evolving technologies are described in a set of three papers, on the futures of CMOS and interconnection technologies by Isaac and Theis, respectively, and on silicon:germanium-based mixed-signal technology by Meyerson. In silicon technology, unlike storage technology, there are no alternates that are expected to play a significant role in less than ten years. Continued progress will require going beyond traditional CMOS scaling, and will require more progress on novel devices, interconnections, and novel circuits. These are discussed with a focus on mixed-signal technologies.

Finally, in the area of displays, as discussed by Wisnieff and Ritsko, advances have also been dramatic. The ability to scale the basic underlying silicon transistor back-planes in an active-matrix liquid crystal panel has led to beautiful displays with resolutions approaching the limits of human visual acuity. Here progress should continue, not just in pixel size scaling, but in other display attributes such as brightness, viewing angle, form factor, and power.

Paul M. Horn
Senior Vice President, Research

Guest Editor